

2008

Generating Ambiguity in the Laboratory

Jack Douglas Stecher

Timothy W. Shields

Chapman University, shields@chapman.edu

John Dickhaut

Chapman University

Follow this and additional works at: http://digitalcommons.chapman.edu/esi_working_papers



Part of the [Econometrics Commons](#), [Economic Theory Commons](#), and the [Other Economics Commons](#)

Recommended Citation

Stecher, J.D., Shields, T.W., & Dickhaut, J. (2008). Generating ambiguity in the laboratory. ESI Working Paper 08-06. Retrieved from http://digitalcommons.chapman.edu/esi_working_papers/143

A peer-reviewed version of this paper was later published as:

Stecher, Jack, T.Shields, and J. Dickhaut. Generating ambiguity in the laboratory. *Management Science* 57.4 (2011): 705-712.

This Article is brought to you for free and open access by the Economic Science Institute at Chapman University Digital Commons. It has been accepted for inclusion in ESI Working Papers by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

Generating Ambiguity in the Laboratory

Comments

Working Paper 08-08

Generating ambiguity in the laboratory*

Jack Douglas Stecher[†]

Timothy Shields[‡]

John Dickhaut[§]

April 2010

Abstract

This article develops a method for drawing samples from which it is impossible to infer any quantile or moment of the underlying distribution. The method provides researchers with a way to give subjects the experience of ambiguity. In any experiment, learning the distribution from experience is impossible for the subjects, essentially because it is impossible for the experimenter. We describe our method mathematically, illustrate it in simulations, and then test it in a laboratory experiment. Our technique does not withhold sampling information, does not assume that the subject is incapable of making statistical inferences, is replicable across experiments, and requires no special apparatus. We compare our method to the techniques used in related experiments that attempt to produce an ambiguous experience for the subjects.

Keywords: ambiguity; Ellsberg; Knightian uncertainty; laboratory experiments; ignorance; vagueness

JEL Classifications: C90; C91; C92; D80; D81

*Thanks to Jonas Andersson, Cris Calude, Monica Capra, Mike Dineen, Miroslav Dudik, Nikolaos Georgantzis, Geoff Gordon, John Hey, Tom Issaevitch, Todd Kaplan, Jack Kareken, Semih Koray, Terje Lensberg, Mark Machina, Charles Noussair, Charles Plott, Remzi Sanver, Teddy Seidenfeld, Erik Sørensen, Bertil Tungodden, Peter Wakker, participants at the Eighth Conference of the Society for the Advancement of Economic Theory, the 38th meeting of the European Mathematical Psychology Group, the 2008 meeting of the Association of Southern European Economic Theorists, and workshop participants at Bilgi, Bilkent, Carnegie Mellon, Instituto de Empresa, University of Haifa, Hebrew University, and University of Zürich. All errors are our own.

[†]Corresponding author. Carnegie Mellon University. E-mail: jstecher@nhh.no.

[‡]Chapman University. E-mail: shields@chapman.edu.

[§]Chapman University.

1 Introduction

This paper develops a method for drawing random numbers from which it is impossible to infer any quantile or moment of the underlying distribution. We mean ‘impossible’ in a strong sense: even the experimenter, with full knowledge of the procedure, cannot make any valid probabilistic inferences, no matter how many data points are observed.

This method provides researchers with a way to give subjects the experience of ambiguity, in the same sense that allowing subjects to flip a coin provides a way to give subjects the experience of probability.¹ This enables experimenters to present ambiguous choices to their subjects without withholding information or finding other ways to proxy for ambiguity: the researcher can study directly the phenomenon that he or she intends to study. The focus can then be entirely on subjects’ behavior (e.g., whether they act as if they face ambiguity, or whether they act like Bayesians and assign priors to any uncertain event, or perhaps whether some other suggested decision heuristic explains behavior particularly well) without any need for concern over whether the data generating process is sufficiently ambiguous. We show mathematically why this is true, illustrate this in simulations, and then test our method in a laboratory experiment.

Our subjects saw 100 histograms of 3,000 realized draws, i.e., a total of 300,000 observations, before we asked them to make any decisions. There is no harm in showing subjects so many draws, because the data generated with our method do not satisfy a law of large numbers. Consequently, our technique can be repeated arbitrarily often. This distinguishes our method from other approaches, which we discuss below in Section 2. And since we do not require any special apparatus, our technique can easily be replicated across experiments.

The main idea behind our method is to set up a data generating process where the cumulative distribution function (cdf) is not always well-defined. This idea goes back to Boole (1854), who

¹Davidson et al. (1957) gave their subject a die, marked on its faces with nonsense syllables that psychology research from the 1920s had shown to have little or no association value. By allowing their subjects to roll the die, Davidson et al. enabled their subjects to reach their own conclusions on what it means to make a decision under risk. We claim here that our method is an analogous way of allowing subjects to experience rather than be told about ambiguity.

argues that the probability that an unknown coin will land Heads should not be considered a definite value like $1/2$, but instead an indefinite value such as $0/0$. Boole clearly has a frequentist interpretation in mind, as the unknown coin would presumably be one on which the observer has seen zero successes on zero tries. Our approach is a slight variation on Boole's idea: we construct a data generating process where the cdf is pointwise equal to $\infty - \infty$, and where the left and right limits of the cdf do not converge. This gives us the undefined order statistics and moments that we need, at least for an objective lack of probability. A lack of subjective probability is ultimately an empirical question, and as we note above, we address this in the laboratory.

The paper proceeds as follows: Section 2 discusses the background and related research. Section 3 presents a description of the technique and shows the results of simulations. Section 4 documents our experimental testing and compares the results with prior work. Section 5 concludes. Appendix A discusses robustness issues. Appendix B provides the instructions for our experiment.

This article develops a method for drawing random numbers from which it is impossible to infer any quantile or moment of the underlying distribution. We mean *impossible* in a strong sense: even the experimenter, with full knowledge of the procedure, cannot make any valid probabilistic inferences, no matter how many data points are observed.

This method provides researchers with a way to give subjects the experience of ambiguity, in the same sense that allowing subjects to flip a coin provides a way to give subjects the experience of probability.² This enables experimenters to present ambiguous choices to their subjects without withholding information or finding other ways to proxy for ambiguity: the researcher can study directly the phenomenon that he or she intends to study. The focus can then be entirely on subjects' behavior (e.g., whether they act as if they face ambiguity, or whether they act like Bayesians and assign priors to any uncertain event, or perhaps whether some other suggested decision heuristic explains behavior particularly well), without any need for concern over whether the data-generating process is sufficiently ambiguous. We show mathematically why this is true, illustrate this in

²Davidson et al. (1957) gave their subject a die, which was marked on its faces with nonsense syllables that psychology research from the 1920s had shown to have little or no association value. By allowing their subjects to roll the die, Davidson et al. enabled their subjects to reach their own conclusions on what it means to make a decision under risk. We claim here that our method is an analogous way of allowing subjects to experience, rather than be told about, ambiguity.

simulations, and then test our method in a laboratory experiment.

Our subjects saw 100 histograms of 3,000 realized draws, that is, a total of 300,000 observations, before we asked them to make any decisions. There is no harm in showing subjects so many draws because the data-generated with our method do not satisfy a law of large numbers. Consequently, our technique can be repeated arbitrarily often. This distinguishes our method from other approaches, which we discuss in Section 2. And because we do not require any special apparatus, our technique can easily be replicated across experiments.

The main idea behind our method is to set up a data-generating process in which the cumulative distribution function (cdf) is not always well-defined. This idea goes back to Boole (1854), who argues that the probability that an unknown coin will land heads should not be considered a definite value like $1/2$, but instead an indefinite value such as $0/0$. Boole clearly has a frequentist interpretation in mind, as the unknown coin would presumably be one on which the observer has seen zero successes on zero tries. Our approach is a slight variation on Boole's idea: we construct a data-generating process in which the cdf is pointwise equal to $\infty - \infty$ and in which the left and right limits of the cdf do not converge. This gives us the undefined order statistics and moments that we need, at least for an objective lack of probability. A lack of subjective probability is ultimately an empirical question, and as we noted earlier, we address this in the laboratory.

The article proceeds as follows: Section 2 discusses the background and related research. Section 3 presents a description of the technique and shows the results of simulations. Section 4 documents our experimental testing and compares the results with prior work. Section 5 concludes. Appendix A discusses robustness issues. Appendix B provides the instructions for our experiment.

2 Background and Related Research

Knight (1921) distinguishes risk from uncertainty as follows: *risk* means that probabilistic information is present, and *uncertainty* means that probabilistic information is absent. A laboratory researcher needs to choose how to convey uncertainty, in Knight's sense, to the subjects in an experiment. We begin this section with a review of how other researchers have implemented decision

making under Knightian uncertainty and then provide general comments on the consequences of different implementation choices. We keep our overview fairly narrow as a comprehensive discussion of this area is available in the work of Wakker (2008).

2.1 Withholding information

Starting with Ellsberg (1961), the tradition in experiments on uncertainty has been to withhold information. This differs from our approach because the uncertainty in the Ellsberg tradition is purely epistemic in nature: an objective probability exists, but the decision maker may be believed to lack sufficient information to form an opinion. Observe that if the experimenter wishes to study repeated decision making, then withholding information becomes difficult as experience will reduce subjects' ignorance. But even if there is no repetition intended, an experimenter loses some control when implementing uncertainty by withholding information.

The way subjects react to withheld information appears to be task-specific. In an early experiment, Chipman (1960) confronts subjects with matchsticks split into heads and stems, placed into various boxes in unknown proportions. After selecting a sample from each box (which, unbeknownst to the subjects, he chooses to be representative), Chipman has his subjects estimate the proportion of heads and stems in each box. He remarks, "One of the most striking features shown by the data is a tendency for individuals to bias unknown probabilities towards one-half" (Chipman 1960, p. 27).

Chipman views this as bias, but Keynes (1921) (in a remark he traces to (Boole 1854)) notes that the issue may simply be one of combinatorics. If a box contains 100 match pieces, each equally likely to be a head or a stem, then there are more compositions of boxes with proportions close to $1/2$ than there are boxes with extreme compositions. In fact, over 99.9% of all compositions would have a proportion of heads between $1/3$ and $2/3$. Subjects who sample from a given box and find that 7 out of 10 sampled pieces are stems could have good reason to believe that the mixture is closer to $1/2$ than to $7/10$. If the subjects treat all possible compositions as equally likely, then before drawing a sample, the subjects would estimate an even mixture of 50 heads and 50 stems to be over 3,000 times more likely than a mixture with 70 stems and 30 heads. The probability of

drawing 7 stems and 3 heads from a box containing 70 stems and 30 heads is slightly more than double that of getting the same draw from a box with a 50-50 mixture. So the subjects would still optimally estimate a probability of close to $1/2$. Indeed, *not* doing so would be an instance of the base rate fallacy (Kahneman and Tversky 1973).

Yet if Chipman's results can be explained by subjects treating all possible compositions of match-boxes as equally likely, the results in other experiments with withheld information cannot. For instance, Einhorn and Hogarth (1986) design their anchor-and-adjust model in part to test a conjecture of Ellsberg's, namely, that subjects prefer uncertainty to risk when facing small probabilities (or appropriately related concepts in the uncertainty tasks). To test for this possibility, Einhorn and Hogarth ask their subjects to choose an urn from which the experimenters would draw a ball. Their risky urn contains 1,000 balls, numbered sequentially. Their uncertain urn also contains 1,000 balls, all numbered (presumably between 1 and 1,000), but Einhorn and Hogarth do not disclose how often each number occurs in this urn. Einhorn and Hogarth then offer their subjects a hypothetical lottery, which awards a prize only if one specific number (say, 672) is drawn.

The results of Einhorn and Hogarth are consistent with their model and are inconsistent with the behavior of Chipman's subjects. If Einhorn and Hogarth's subjects had regarded each ball in the uncertain urn as having equal likelihood of being labeled with any integer from 1 to 1,000, then the subjects should have been indifferent between the two urns. But what if, instead, Einhorn and Hogarth's subjects viewed all ratios as equally likely? It is easy to imagine subjects who know that the winning number is printed on anywhere from $0/1000$, $1/1000$, \dots , $1000/1000$ of the balls in the urn. Under this interpretation, the uncertain urn would be expected to give the prize with probability $1/2$ and hence would have 500 times the expected value of the risky urn.

To recap, there are at least two difficulties that withholding information raises from an implementation standpoint. The first is that it is hard to keep subjects ignorant as they can always choose to act like frequentists. So any task that involves repeated decision making would require new sources of uncertainty at each stage, which subjects could not infer from their previous observations. The second is deeper: even in a single-task experiment, a subject facing withheld information may interpret the information in a task-specific way. Subjects in Chipman's experiments act consistently with viewing all possible compositions as equally likely, whereas those in Einhorn and Hogarth's

experiment act consistently with viewing all possible ratios as equally likely. Unless the experimenter has a way to control which of these interpretations would be more natural in a given task, he or she may generate results that are not straightforward to evaluate.

2.2 Complicated and compound lotteries

Rather than withhold information, several studies have tried to draw from a distribution that the subjects could see, yet might still view as difficult to specify. These are similar to our work in spirit in that they attempt to separate ambiguity from withholding information.

Hayashi and Wada (2006), in a working paper version of their subsequent paper (Hayashi and Wada (2009)), let their subjects generate probabilities. One group of subjects chose mixtures of balls in urns, after which another group made decisions in Ellsberg-style tasks. For the final version of the paper, however, Hayashi and Wada abandon this approach, as human subjects are too systematic. The subjects deciding on mixtures overwhelmingly selected either symmetric distributions, in which all colors of balls were equally likely, or highly asymmetric distributions, where at least one color was excluded.

In the final version of their paper, Hayashi and Wada (2009) use lotteries that draw from a distribution that seems difficult to estimate. As a source of uncertainty, Hayashi and Wada use the probabilities of outcomes of snakes and ladders games, which they view as too complex for subjects to estimate. Unlike methods based on withholding information, this approach allows the researchers to disclose everything about their procedures to their subjects. It should be noted, however, that the distribution of outcomes of snakes and ladders games is actually fairly simple (for details, see Berlekamp et al. 1982, Althoen et al. 1993). Variations on this technique may find a more complicated distribution, but any distribution satisfying a law of large numbers will still be difficult to use in studies of repeated decision making. In other words, the fact that a well-behaved distribution exists in the Hayashi and Wada approach limits the studies in which an experimenter can apply it.

Hey et al. (2007) use a British bingo blower, which is a device similar to those used to select Ping Pong balls in state-run lotteries. As the balls are constantly in motion, it is difficult for subjects to form an accurate picture of how many balls of each color are in the machine. As Hey et al.

note, there is no dispute that there is a fixed objective distribution. So this approach has the same inherent limitation as the Hayashi and Wada approach.

An added difficulty of the Hey et al. approach is that it requires the researchers to withhold information from subjects, as in Ellsberg. The properties of machines that pick Ping Pong balls for state-run lotteries are discussed by Stern and Cover (1989). They show that the empirical distribution of balls drawn from such machines passes every statistical test of being uniform. Thus the ambiguity in Hey et al. must be in the first stage of the lottery, when the subjects must estimate the proportions of balls in the machine.

An advantage of the Hey et al. and Hayashi and Wada methods is that they each enable experimenters to show their subjects directly what procedures are in use. This contrasts with the Ellsberg method of withholding information, in which the researcher tells the subjects what they know (and, perhaps somewhat strangely, what they do not know). However, a drawback of the Hey et al. and Hayashi and Wada approaches is that they have room for comparative ignorance. Subjects may have more or less experience with snakes and ladders, a game that is essentially a Markov chain, and may therefore differ in knowledge of what outcomes are likely, even if they do not know the underlying mathematics. Similarly, one subject may have a relatively good feel for how to estimate the proportion of blue balls swirling through a bingo blower or similar device, whereas another may be terrible at guessing how many blue balls are floating through some sort of light blue mass of balls in motion. The latter subject would be relatively more ignorant than the first about the distribution of balls, and quite possibly, both subjects could have some idea about how well or poorly they make these sorts of estimates. Because comparative ignorance is known to affect how subjects make decisions under uncertainty (see, e.g., Fox and Tversky 1995), this restricts the types of questions one can study with these techniques.

A related approach uses compound lotteries. If subjects lack the skill to reduce compound lotteries to simple ones in the time allotment of an experiment, they might rely on computationally less burdensome techniques for finding lower and upper bounds on the payoffs of a lottery. Arló-Costa and Helzner (2009) test decision making with compound lotteries and compare it to subjects' decisions when facing withheld information and when facing risk. They generate distributions of balls in an Ellsberg-type urn using a uniform distribution over possible ratios of red balls to black

balls. They explain this procedure to their subjects and illustrate it with a roulette wheel example. In a separate task, their subjects price a lottery with the standard Ellsberg two-color description.

The pricing tasks in Arló-Costa and Helzner shed some light on how subjects' reactions to compound lotteries compare with their reactions to withheld information. Their subjects are asked to price simple lotteries, compound lotteries in which each winning ratio $0/100$, $1/100$, \dots , $100/100$ is equally likely, and the Ellsberg two-color lottery. Their subjects consistently price the simple lottery highest, at roughly double the value of the compound lottery, which in turn is roughly double the value of the Ellsberg lottery. Thus the subjects' behavior is as follows:

$$\frac{\text{price of ambiguous}}{\text{price of compound}} \approx \frac{\text{price of compound}}{\text{price of simple}}.$$

This suggests that compound lotteries are seen as an intermediate case between ambiguity and risk and not as a proxy for ambiguity itself. One caveat is that the pricing differences between the compound and ambiguous lotteries become small when they compare subjects whose first decision is about the compound lottery with those whose first decision is about the ambiguous lottery. Even so, when the ambiguous lottery is presented first, the subjects give a higher price to the compound lottery, and when the compound lottery is presented first, the subjects give a lower price to the ambiguous lottery. So though there is suggestive evidence that subjects react similarly to compound lotteries as they do to ambiguous lotteries, this evidence is restricted to comparison across experiments. These results confirm earlier findings by Yates and Zukowski (1976) and Chow and Sarin (2002), whose subjects likewise react to compound lotteries as if facing an intermediate case.

2.3 Refining the notion of ambiguity

An advantage of the technique we propose, in contrast to those discussed in Sections 2.1 and 2.2, is that our method has no well-defined (objective) cumulative distribution function. This means that there is no need for an experimenter to conceal the procedure from one's subjects and that there is no difficulty in repeating or replicating an experiment.

To clarify this distinction, we introduce distinctions among terms that are often used interchangeably in the literature. For the remainder of this article, *Knightian uncertainty* will refer to the

general setting in which probabilistic information is absent, irrespective of the reason. We call the absence of objective probability *ambiguity* and will henceforth restrict our use of this term to settings in which a cdf is objectively ill-defined. If an objective distribution is present, but the subjects may lack the required information to find it or even to form well-defined priors, then we will say that the subjects are making decisions under *ignorance*. This was the term Hurwicz (1951) introduced in early work on such situations. Finally, if a task is given where multiple interpretations are possible, as in the Chipman and Einhorn and Hogarth studies mentioned in Section 2.1, we will refer to the task as one of decision making under *vagueness*. This matches use in the philosophy literature and is also the term that Arló-Costa and Helzner (2009) use to describe the Ellsberg-type urn, in contrast to their compound lottery (which may involve ignorance, but not vagueness). We can now restate our claim: our technique generates ambiguity in this narrow sense, and doing so avoids some of the disadvantages of prior approaches to Knightian uncertainty, which have relied on vagueness or ignorance. Before giving more details on our approach, we elaborate on the underlying theory and historical theoretical discussion.

The canonical example of ambiguity in our sense is a nonmeasurable event; Binmore (1995, 2007a,b, 2009) pursues this idea in a sequence of papers and a recent book. Intuitively, a lower bound on the probability of a nonmeasurable event is given by an inner measure, and an upper bound is given by an outer measure. If these converge, a unique probability exists, and if the outer and inner measures do not converge, probabilities are nonunique.

The idea we pursue is closely related: a data-generating process is ambiguous at any point where its cdf is undefined and where its left limit and right limit do not converge. If x is a point in the support of the cdf $F(\cdot)$ of some data-generating process, then

$$F(x) = Pr(\{\omega \in \Omega | X(\omega) \leq x\}),$$

where Ω is the set of states of the world and $X(\cdot)$ is the observed random variable. If $F(x) = \infty - \infty$, we might try to estimate the value of F as we approach x from below and from above:

$$F(x)^- := \lim_{\hat{x} \nearrow x} F(\hat{x}) \quad F(x)^+ := \lim_{\hat{x} \searrow x} F(\hat{x}).$$

If these are well-defined, then we can view $F(x)^-$ as a lower bound on the probability of getting a draw less than or equal to x and $F(x)^+$ as an upper bound on this probability. If $F(x)^-$ does not

exist, we can bound $F(x)$ below by picking the largest $y \leq x$ such that $F(y)$ exists, or in the worst case at 0. Similarly, if $F^+(x)$ does not exist, we can bound $F(x)$ above by picking the smallest $z \geq x$ such that $F(z)$ exists, or in the worst case at 1.

Other authors have pursued this idea in slightly different frameworks (see, e.g., Suppes 1974, Good 1983, Fishburn 1993, Sutton 2006). This idea is consistent with the views of Kolmogorov (1983), who writes that probabilities are ill-defined in “most” cases and that stochastic analysis is limited to the study of the special cases in which probabilities are uniquely defined. Papamarcou and Fine (1991) and Fierens et al. (2009) provide examples of physical phenomena that cannot be characterized by unique probabilities, hence where this approach would fit.³

Discussions of vagueness go back at least to the 19th century. Jevons (1874) illustrates vagueness, considering the statement “A Platythliptic Coefficient is positive” and asking what reasonable probability one can assign to the likelihood that this statement is true.⁴

The problem of vagueness arises with withholding information, as the results of the Chipman and Einhorn and Hogarth experiments, discussed in Section 2.1, suggest. These results were anticipated by Keynes (1921), who in turn traces his argument back to Boole (1854). Keynes considers an urn containing only black and white balls in unknown proportion and asks what it means for someone to consider all possible urns equally likely.

To illustrate, imagine an urn containing three balls, each of which is either black or white, as Keynes suggests. The possible proportions of black balls are 0/3, 1/3, 2/3, and 3/3. One interpretation of

³Girón and Ríos (1980) propose an extension of Bayesian decision making to environments such as these, which they call “quasi-Bayesian” decision theory. Some other well-known extensions are the models of Gilboa and Schmeidler (1989, 1993), Klibanoff et al. (2005), Maccheroni et al. (2006), and Hanany and Klibanoff (2007). Seidenfeld and Wasserman (1993) discuss a difficulty with a Bayesian-inspired approach: updated nonunique beliefs may be a superset of the priors. Earlier approaches, such as the minimax regret rule that Savage (1972) proposes, maximin rules, and a principle of insufficient reason, are axiomatized by Milnor (1954).

⁴He argues that the only sensible choice is 1/2; in contrast, Boole takes position that the correct value is an indefinite number like 0/0. Keynes (1921) sides with Boole, illustrating the issue by rephrasing Jevons’s question. He asks what probability Jevons would assign to the proposition “A Platythliptic Coefficient is a perfect cube” or to “A Platythliptic Coefficient is allogeneous.” These are just as vague as Jevons’s suggestion, so should they also have probability 1/2 of being true?

equal likelihood is that each of these ratios is as likely as any of the others. On the other hand, each of the balls can be one of two colors, so there are $2^3 = 8$ possible compositions of the urn:

WWW, WWB, WBW, BWW, WBB, BWB, BBW, BBB.

Each composition may be deemed equally likely, on the grounds that there is no reason any ball should favor being black or white. The first composition has no black balls. The next three compositions each have one black ball. The next three compositions each have two black balls, and the last has three black balls. Thus we are led to two different possible distributions of balls in the urn, even though, in both cases, we are acting as if all possibilities are “equally likely.”

As with ignorance, vagueness decreases with experience. Thus one reason for distinguishing ambiguous probabilities from ignorance or vagueness is repeatability of a task within an experiment. Another is replicability across experiments: different populations may vary in how consistently individuals resolve semantic vagueness or in how they update their beliefs when they start out ignorant. More fundamentally, separating ambiguity, ignorance, and vagueness is useful for investigating which forces drive the results in well-studied models of decision making under uncertainty.

We now present our technique for drawing ambiguously distributed numbers. It has no unique distribution to estimate and does not satisfy a law of large numbers. Hence there is no need to hide anything from the subjects, no room for comparative ignorance, and no issue of vagueness.

3 Description of the Technique

3.1 The basic technique

We present our basic technique here. For now, we focus attention on the case in which an ambiguous sample is needed but in which the subjects do not observe each draw sequentially. It turns out that adjusting for sequential observations only requires repeated application of our basic procedure. After introducing our basic procedure, we illustrate it with simulations and then turn to the case of sequential observations. We end this section with a discussion of practical implementation issues.

Our main idea is to work with a second-order, nonstationary Cauchy process. Very roughly speak-

ing, we begin with Cauchy distributed random variables, which are known to present difficulties for inference, and then destroy the remaining properties that could be of use to an idealized statistician. In the end, we have a process that we can fully characterize, yet on which no inference about any order statistic is possible. This is because the sample order statistics are divergent. These two features are the essence of what makes our technique work. First, because we can fully characterize our procedure, we can draw from it as often as we like, making it replicable across experiments. Second, every order statistic being unpredictable means, essentially, that the lower bound on the probability of a draw being in a given set will not converge to the upper bound.

The Cauchy distribution is a convenient point of departure because it has no moments about any number. This property does not make probabilities ill-defined, but it is a start. The definitions of the Cauchy cdf and density are as follow:

$$F(x) = \frac{1}{\pi} \arctan\left(\frac{x-x_0}{\gamma}\right) + \frac{1}{2} \quad f(x) = \frac{1}{\pi\gamma \left[1 + \left(\frac{x-x_0}{\gamma}\right)^2\right]}.$$

The parameter x_0 is called the *location* of the distribution, and the parameter γ is called the *scale*. These fully characterize the distribution. Subtracting x_0 from a draw and dividing by γ is a standardization, analogous to calculating a z -score. The arctan function ranges from $-\pi/2$ to $\pi/2$ and increases in its argument. Dividing by π and adding $1/2$ gives a function that ranges from 0 to 1 and is strictly increasing, hence an atomless cdf on \mathbb{R} . The derivative of $\arctan x$ is $1/(1+x^2)$, so the probability density function is easily derived from the chain rule. Figure 1 compares the standard Cauchy and standard normal densities. The striking feature of the Cauchy density is its leptokurtosis.

It is known that, for a Cauchy distributed random variable X and any $\hat{x} \in \mathbb{R}$,

$$E[X|X \leq \hat{x}] = -\infty \text{ and } E[X|X > \hat{x}] = +\infty.$$

Because one can make $\infty - \infty$ equal anything, there is no mean of a Cauchy distributed random variable. The second moment about any number, hence any replacement for variance, of a Cauchy distributed random variable is infinite. More generally, all the odd-numbered moments are divergent, and all the even-numbered moments are infinite. As finite means and variances underlie laws of large numbers and central limit theorems, it is unsurprising that neither holds for a Cauchy random variable. This alone makes statistical inference difficult in the Cauchy case.




Figure 1: Cauchy density (solid) compared with normal density (dashed).

Despite these properties, the Cauchy distribution will not serve our purpose without further modification. Barnett (1966) shows that the sample order statistics are good estimators of the Cauchy distribution's theoretical quantiles. So while many common descriptive statistics are uninformative, the cumulative distribution function, and hence whatever probabilities one wants, can be estimated. For instance, the sample median is a good estimator of the location x_0 . With any additional sample quantile, one can estimate γ , and hence $F(\cdot)$. We need something less well-behaved.

Because a Cauchy random variable has no mean, we define the following process X_t :

$$X_0 \sim C[0, 1] \quad X_{t+1} \sim C[x_t, 1],$$

where we adopt the convention of using lowercase letters to indicate realizations and capital letters to indicate random variables. In other words, we draw our first observation from a Cauchy distribution with a location of 0 and a scale of 1 (called a *standard* Cauchy). For each subsequent observation, we draw from a Cauchy distribution with a location equal to the previous realization. The expected location of any future draw is the mean of a Cauchy random variable, which does not exist.

The preceding procedure still misses our goal. The reason is that drawing from a Cauchy distribution with location $x_0 = x$ and scale $\gamma = 1$ is equivalent to drawing from a standard Cauchy and then adding x . Pitman and Williams (1967) show that many functions, including sums, of Cauchy random variables are surprisingly well-behaved. For the process X_t , the n th draw is Cauchy, with $x_0 = 0$ and $\gamma = n$. The sample order statistics remain informative.

However, we are on the right track. A second-order Cauchy process can make both the scale γ and the location x_0 depend on realizations. This enables us to define a process that is nonstationary, nonergodic, and does not meet the conditions of Pitman and Williams. The resulting process has truly horrendous sample order statistics, as we would hope. There are two caveats: we must make certain that γ is always positive, and we must prevent γ from growing too rapidly. To see this latter point, note that a large value of γ increases the likelihood of extreme draws. This increases γ further, causing the process to generate values that may create overflows in software applications.

Accordingly, our basic procedure works as follows:

1. Draw $Z_0 \sim C[0, 1]$.
2. Draw $Z_1 \sim C[z_0, 1]$.
3. Let $\phi, \psi \in (0, 1)$ with ϕ, ψ both small. For $n \geq 2$, draw $Z_n \sim C[z_{n-1}, \phi \cdot |z_{n-2}| + \psi]$.

The parameters ϕ and ψ ensure that the scale parameter is nonzero. Picking both close to zero slows the tendency of the resulting data to diverge toward $\pm\infty$.

The preceding procedure is easily modified for drawing ambiguously distributed numbers with finite or bounded support. We restrict attention to the case of finite support here as the technique for generating bounded support is nearly identical. The main trick is to use a latent second-order Cauchy process. Because this latent process is nonergodic, its behavior when mapped onto a finite set is unstable, which is what generates the desired properties.

Let $k \in \mathbb{N}, k \geq 2$ be the number of possible outcomes of the experiment, such as the number of colors of balls in an Ellsberg-type urn. The procedure is as follows:

- Generate a sequence $\{z_n\}$ as a second-order Cauchy process, as described earlier.
- Define a new sequence $\{z'_n\}$ by setting z'_i to the greatest integer no larger than z_i :

$$z'_i := \text{Floor}(z_i).$$

- Now define a third sequence $\{y_n\}$ by taking each $z'_i \bmod k$ and adding 1:

$$y_i := (z'_i \bmod k) + 1 = (\text{Floor}(z_n) \bmod k) + 1.$$

Adding 1 to the result just labels the outcomes $1, \dots, k$.

For example, suppose $z_i = -23.7894$ and $k = 10$. The floor of z_i is -24 . Now, take the remainder of dividing by 10. Because -30 is the largest multiple of 10 that is no larger than -24 , the remainder is $-24 - (-30) = 6$. Add 1 to get $y_i = 7$. If z_i has been positive 23.7894 , we would have obtained $23 \bmod 10 + 1 = 4$.

3.2 Simulations

We use two approaches to illustrate our technique. First, we show the results of simulations using pseudo-random number generators. This enables us to demonstrate the divergence of the sample order statistics. Second, we use our technique in a laboratory experiment, based on the Ellsberg two-color urn. We describe the simulations here and the experiment and its results in Section 4. In Appendix A, we provide simulations using physical randomization devices.

For the simulations, we drew 10,000 ambiguously distributed random variables, both with unbounded support and with support in $\{1, \dots, 10\}$. We then replicated the procedure 100 times. We chose $\phi = 0.001$ and $\psi = 0.0001$.

For each unbounded sequence of 10,000 draws, we computed several order statistics and then compared how the sample order statistics behaved over the 100 replications. The range of values for each sample order statistic, along with the median and mean values, are shown in Table 1.

None of the sample order statistics converged after $10,000 \cdot 100 = 1$ million draws. For example, the minimum was below -3 million in one replication and was strictly positive (just over 9) in another. The maximum ranged from below -2 to above 18 trillion. So at least one replication had 10,000 negative draws, and at least one other had 10,000 positive draws. The most stable sample order statistic was the 10 %ile, where the highest value exceeded the lowest by more than 54,000.

Table 1: Sample order statistics, 100 replications of 10,000 draws

<i>Order Statistic</i>	<i>Minimum</i>	<i>Median</i>	<i>Mean</i>	<i>Maximum</i>
Minimum	$-3.32017 \cdot 10^6$	-66.2505	-39718.8	9.20083
1 %ile	-562905	-51.9015	-11119.7	11.292
2.5 %ile	-145903	-40.9968	-4914.81	12.016
5 %ile	-96887	-29.9379	-3774.48	23.3314
10 %ile	-54679.7	-10.7992	-1172.47	36.0455
Lower Quartile	-33207.8	-1.45783	53.3294	47618.7
Median	-5348.3	0.0298731	2963.11	296289
Upper Quartile	-1196.73	1.70167	$1.66938 \cdot 10^7$	$1.66926 \cdot 10^9$
90 %ile	-371.732	7.91172	$7.09551 \cdot 10^7$	$7.09474 \cdot 10^9$
95 %ile	-342.251	19.2627	$1.06024 \cdot 10^8$	$1.06006 \cdot 10^{10}$
97.5 %ile	-27.9362	30.8853	$1.20062 \cdot 10^8$	$1.20025 \cdot 10^{10}$
99 %ile	-24.6536	39.7169	$1.33105 \cdot 10^8$	$1.33064 \cdot 10^{10}$
Maximum	-2.16994	46.6691	$1.88546 \cdot 10^8$	$1.88503 \cdot 10^{10}$

For finite support, we converted the simulated draws to the range $\{1, \dots, 10\}$. Table 2 summarizes the first ten replications. Again, each replication consisted of 10,000 draws.

Table 2: Sample histogram, finite support, 10 replications

<i>Replication</i>	1	2	3	4	5	6	7	8	9	10
I	28	499	795	783	1,520	5,593	512	36	0	234
II	1,016	837	483	573	244	46	768	2,477	1,467	2,089
III	1,521	639	251	575	902	827	1,042	2,323	974	946
IV	0	0	0	0	0	0	54	2,084	6,833	1,029
V	2,220	236	580	603	632	912	758	550	378	3,131
VI	666	712	751	842	1,725	1,641	1,191	707	845	920
VII	17	119	187	483	7,952	754	292	31	102	63
VIII	315	360	757	5,901	495	723	360	226	569	294
IX	2,261	1,034	378	90	77	221	92	152	79	5,616
X	1,795	1,658	861	298	106	302	303	32	1,920	2,725

In the first replication, the most common draw was 6, making up 55.93% of the realizations, while 9 was not drawn at all. In the second replication, 14.67% of the draws (just over $1/7$) were equal to 9, whereas only 0.46% were equal to 6. These were not continuations of the same process: for each replication, we restarted. So the procedure was identical, but the results differed wildly.

The results did not stabilize over the full set of 100 replications. For each outcome in $\{1, \dots, 10\}$, there was at least one replication in which it was never drawn, and there was at least one replication when it made up at least 28.5% (roughly $2/7$) of the draws. Eight of the 10 outcomes had at least

one replication where it made up more than half of the draws. The most extreme replication had 88.92% of the draws equal to 1, 11.08% equal to 2, and no other values realized. The most symmetric case had each outcome occur at least 8.8% of the time and at most 11.47% of the time. Even this case is inconsistent with the data coming from a uniform distribution.

3.3 Sequential observations

The basic procedure just described suffices for experiments in which subjects will not see the data sequentially, such as the experiment we present in Section 4. In other experiments, however, the subjects may observe draws one by one. Because the current draw depends on the realizations of the previous two draws, it is useful to garble the order in which the draws arrive.

Figures 2 and 3 illustrate this issue. Each figure displays 10,000 draws from the second-order Cauchy process in sequence. In Figure 2, the process is fairly stable, with values staying close to zero for the first 1,000 draws. Volatility increases between draws 1,000 and 4,500, with a general downward drift toward around -20 . The next 2,000 draws show a sharp increase in volatility, with the process jumping up to a value of around 60 near draw 6,000. The values then drop, and the process becomes less volatile for the remainder of the time series. By contrast, the path in Figure 3 shows a much wider range of values. The first 5,000 draws appear almost flat, but this is really an artifact of the wild volatility that arises in the time series shortly after the 5,000th draw. From there, the process ranges from $-4,000$ to just above $-1,000$.

The Cauchy process has periods of stability occur until an extreme draw arrives; this happens often because of the leptokurtosis of the Cauchy distribution. Each extreme draw gets embedded (two draws later) into the scale parameter. This makes subsequent extreme draws more likely, causing the process to enter an unstable period. The length of these periods depends on the scale parameter, which itself is random and has no mean. This makes the period length unpredictable.

The data nevertheless are serially dependent and the subjects could take advantage of whether the process is in a calm or unstable period. One remedy is to shuffle: randomly permute the original data, then present the data in this garbled order. Scheinkman and LeBaron (1989) use this approach when testing for serial dependence in stock prices. They create simulated time series




Figure 2: Simulation of 10,000 draws.

of financial returns by shuffling the real data, and then observe characteristics that are common to the real data but unobserved in the shuffled data.

For some experiments, shuffling may be good enough. But shuffling does not fully destroy serial dependence in a data series, as has been known at least since DeGroot and Goel (1980). The problem is the same combinatorial issue mentioned earlier: not all gaps between draws that were originally adjacent are equally likely. For example, imagine a sequence of three draws, x_1, x_2, x_3 , from some process where x_1 and x_2 are not independent. There are six possible permutations:

$$x_1, x_2, x_3 \quad x_2, x_1, x_3 \quad x_3, x_1, x_2 \quad x_3, x_2, x_1 \quad x_1, x_3, x_2 \quad x_2, x_3, x_1$$

In the first permutations, x_1 and x_2 remain adjacent. So while shuffling makes each permutation equally likely, it does not fully destroy the information present in the time series.

To get around this difficulty, one can use the second-order Cauchy process multiple times. First,




Figure 3: Another simulation of 10,000 draws.

generate a sample z_1, z_2, \dots, z_{n^*} of at least the required size n , but make the actual sample size n^* a stopping time. For instance, one might make the sample size equal to

$$n^* = \max\{n, \text{Sample size necessary to obtain first draw above } 1000\}.$$

Doing this assures that each data set of length n is drawn from a different distribution. Once an initial sample z_1, z_2, \dots, z_{n^*} is drawn, generate a second Cauchy process of length n , taking values in $\{1, \dots, n^*\}$. These will serve as indices. Call this process j_1, j_2, \dots, j_n .

Now use the indices as if bootstrapping the original sample: choose a subsample z'_1, \dots, z'_n from z_1, \dots, z_{n^*} by setting $z'_k = z_{j_k}$. The result is an ambiguous sample, generated from a path of ambiguous length, ordered in an ambiguously defined way.

3.4 Practical comments on choosing coefficients

Our technique balances two opposing forces. The properties of the Cauchy distribution make extreme draws a common occurrence. However, the coefficients ϕ and ψ on the scale parameter γ , if sufficiently small, have a damping effect, making a large number of draws stay close together. The purpose of the additive constant ψ is simply to avoid division by 0, so we concentrate on the effects of ϕ . As ϕ increases, more extreme draws arrive more often. To illustrate, we repeated the simulation of Section 3.2, but increased ϕ from 0.001 to 0.1. We kept $\psi = 0.0001$ as before.

Because of this reduced damping, the behavior of the draws was more dramatic than in our earlier simulations. Table 3 shows the range of the mean and median over the 100 replications:

Table 3: Sample mean and median, $\phi = 0.1$.

<i>Statistic</i>	<i>Minimum</i>	<i>Median</i>	<i>Maximum</i>
Median	$-1.6 \cdot 10^{17}$	0	$1.1 \cdot 10^{19}$
Mean	$-2.3 \cdot 10^{122}$	$-2.5 \cdot 10^{47}$	$5.1 \cdot 10^{110}$

All the other order statistics are similarly erratic. As an additional test, we computed 100 more replications of 10,000 draws, this time increasing ϕ to 0.2. The range of all of the order statistics widened further, with realizations reaching absolute values on the order of 10^{294} .

Data of these absolute magnitudes may not be stored as an integer within a software application. To demonstrate this issue, we took the data series that we had generated by setting $\phi = 0.1$ and converted it to integers from 0 to 9. In all 100 replications, every even number occurred with higher frequency than all of the odd numbers. When we used the data series we had generated by setting $\phi = 0.2$, this pattern became even more pronounced.⁵ Overall, setting the damping parameter to 0.1 or 0.2 does increase the magnitude of the realizations in the unbounded data. However, Section 3.2 demonstrated that instability is already attainable with $\phi = 0.001$, without risking overflows or making the conversion to integers questionable.

Conversely, reducing ϕ makes the frequency of extreme values relatively low. To illustrate this, we drew another 100 replications of 10,000 ambiguously distributed numbers, this time setting $\phi = 0.000001$. We then converted to the case of finite support, choosing support $\{0, \dots, 9\}$. Among

⁵The simulations were done in *Mathematica*.

the 100 replications, there were 82 where at least 9,999 out of 10,000 draws all ended up as the same digit (though the specific digit varied across replications). There were another six replications where two digits accounted for all realizations, and eight more where three digits accounted for all realizations. The remaining four replications included one where four digits accounted for all draws, one where six digits accounted for all draws, and two where all ten digits had some realizations. In one of these last two cases, each digit was realized at least 517 times and at most 1532 times; the other case was slightly less symmetric but similar.

There is evidently a trade-off between a low damping parameter, which keeps the realized values from getting too extreme and causing overflows, and a larger one, which keeps the realized values from being concentrated on a narrow range of outcomes. Even with an extremely low damping parameter, there was a dramatic difference across replications, reinforcing the idea that the underlying distribution of the data is ambiguous.

4 The Technique in Practice: A Laboratory Experiment

To show how our technique works in practice, we conducted laboratory experiments based on the Ellsberg two-color urn. We begin by giving some background on the experimental procedures.

We first independently asked each subject to choose between gambles presented as computerized penny flips and then asked each subject to price each penny flip. We did not disclose individual decisions to other subjects, and each subject's payoff was not affected by the other subjects. Prior to eliciting any choice, we presented the entire experiment, payoffs, and tasks within the instruction set. Each session lasted less than one hour and consisted of the following steps:

1. The experimenter read the instructions aloud.
2. The subjects took a quiz on the experimental tasks.
3. The computer notified each subject of whether heads or tails paid the higher amount.
4. The computer asked the subjects which penny they preferred to flip.

5. The computer asked the subjects to input an asking price for each of two pennies.
6. One of a subject's three elicited choices (the choice, price for the first penny, and price for the second penny) was randomly selected for payment.
7. The computer privately revealed the outcome and payoff to each subject.
8. The subjects privately received payment and signed a compensation form.

We ran multiple sessions of the experiment at a university in the western United States between May and June 2009. In total, we recruited 60 subjects from a standard subject pool, consisting primarily of undergraduate university students. The subjects performed the tasks anonymously over a local computer network. We programmed and conducted the experiment using z-Tree (Fischbacher 2007). The computers were placed within individual cubicles in such a way that each subject could only view his or her own computer screen.

The same staff member facilitated all sessions of the experiment. The staff member distributed printed instructions to the subjects and then read the instructions aloud while the subjects followed along with their own copies. We show the instructions in Appendix B. The instructions included 100 histograms for each penny, each consisting of 3,000 penny flips (see Figures 4 and 5).

In discussing the setting and results, we refer to the gambles as the Risky Penny and the Ambiguous Penny, whereas in the experiment, we labeled the gambles as Penny M and Penny C. Within the experimental instructions, the order of graph collections was randomized over subjects. The Risky Penny represented a Bernoulli distribution with a heads probability of $1/2$, whereas the Ambiguous Penny represented the technique described in Section 3.1, with parameters $\phi = 0.0001$, $\psi = 0.0001$, and mapped to a discrete outcome of heads or tails. Overall, the Risky Penny graphs had 150,262 heads and 149,738 tails, whereas the Ambiguous Penny graphs had 141,334 heads and 158,666 tails. All bar graphs for the Risky Penny displayed approximately half heads or tails. Of the 100 bar graphs for the Ambiguous Penny, 17 displayed all heads or all tails.

Subjects took a quiz on the experimental tasks, price, and payoff mechanisms. An experimenter or one of the authors privately reviewed each subject's answers. On average, the subjects marked one incorrect answer among the eight questions.




Figure 4: Collection of bar charts of the Risky Penny

A collection of bar charts of the Risky Penny, where each chart reports the results of 3,000 flips of the penny. There are 100 such charts, each showing a red bar, denoting the number of times the penny landed heads out of 3,000 flips, and a blue bar, denoting the number of times the penny landed tails out of 3,000 flips.




Figure 5: Collection of bar charts of the Ambiguous Penny

A collection of bar charts of the Ambiguous Penny, where each chart reports the results of 3,000 flips of the penny. There are 100 such charts, each showing a red bar, denoting the number of times the penny landed heads out of 3,000 flips, and a blue bar, denoting the number of times the penny landed tails out of 3,000 flips.

The computer notified each subject on screen if heads paid \$30 and tails paid \$10, or if tails paid \$30 and heads paid \$10. We randomized these two payoff mappings over the subjects.

We elicited three choices from each subject. First, we asked each subject which penny he or she preferred to flip. We randomly varied the presentation order, showing the choice with the Risky Penny on the right and the Ambiguous Penny on the left, or vice versa. After the subjects made the pairwise choice, we asked each subject to input an asking price between \$10 and \$30 for each penny. Again, we randomized the solicitation order over the subjects, first eliciting the price for the Risky Penny and then the Ambiguous Penny, or vice versa.

We randomly selected one of the subject's three elicited choices for payment. If we selected one of the pricing choices, we implemented the Becker-DeGroot-Marschak pricing mechanism (henceforth BDM, due to Becker et al. 1964), in which we randomly drew a buying price between \$10 and

\$30. If the buying price exceeded the asking price, the subject sold the penny at the buying price; otherwise, the subject kept and played the gamble. Lastly, if the penny was not sold, the computer flipped it and calculated the payment. We privately paid the subjects and had them sign a compensation form.

4.1 Experimental results

We found no significant differences among the five sessions and hence focus here on the elicited choices of all 60 subjects and report the combined results. Before assessing the subjects' preferences, we examine effects of order on the choice between the two gambles and on price as well as the effect of heads or tails payment on the coin choice.

The decision to select the coin presented on the left versus the right shows no significant correlation, using the Spearman rank-order correlation coefficient ($n = 60$, $p \leq .33$). To test whether the presentation order of the pennies affected the pricing decision, we compare the price of the first choice presented to that of the second choice. We fail to find a difference using the Wilcoxon matched-pair signed-rank test ($W = 145$, $n_{s/r} = 38$, $p \leq .30$). Lastly, to test whether heads paying \$30 and tails \$10 versus the reverse affected the choice of pennies, we use the Spearman rank-order correlation coefficient and fail to find a correlation ($n = 60$, $p \leq 1$).

Given the pairwise choice between pennies, 36 of 60 subjects chose the Risky Penny over the Ambiguous Penny (see Table 4). This result is consistent with prior research documenting ambiguity aversion. We use the solicited prices to further determine if the subjects viewed the two pennies as statistically equivalent.

The subjects did not en masse price the two pennies equally: 38 of 60 subjects priced one penny higher than the other. Comparing the choice and prices input allows us to infer subjects' preferences. The largest group seemed to have a distinct preference for one penny over the other: 24 of 60 subjects priced the penny they chose higher than the other. Among these 24, 15 preferred the Risky Penny, whereas nine preferred the Ambiguous Penny. The second largest group, 22 of 60 subjects, appeared indifferent between the pennies, pricing both pennies exactly equally. Of those 22 subjects, 14 chose the Risky Penny in the pairwise exercise, compared with 8 who chose the

Ambiguous Penny. Lastly, 14 of 60 subjects' choices were inconsistent, pricing the chosen penny lower than the foregone penny. Our design does not allow us to further expound on these subjects' preferences.

Table 4: Subjects composite choices.

	Priced Risky Penny Higher	Priced Both Pennies Same	Priced Ambiguous Penny Higher	Totals
Chose Risky	15	14	7	36
Chose Ambiguous	7	8	9	24
Totals	22	22	16	60

Thus we had 46 subjects who did not reverse preferences, among whom just over half (24) showed a clear preference, with the other 22 indifferent. A 95% confidence interval over those who did not reverse preferences gives between 37.12% and 66.86% of subjects with a strict preference. If we include the subjects who reversed preferences and assume that we observed the true proportion of preference reversals, then with 95% confidence, more than 26.8% must have had a strict preference, and fewer than 49.9% could have been indifferent. Thus, if the null hypothesis is that subjects view the Risky Penny and Ambiguous Penny as equivalent, then we strongly reject the null.

4.2 Comparisons to existing experimental results

Our results are broadly in keeping with the body of research on ambiguity, although we find some differences with each individual study. We list several other experimental studies on ambiguity in Tables 5 and 6. We base our comparisons on the observations shown in Table 4. Some of the other studies differ from ours in that either (1) they do not include a means of measuring indifference or (2) they do not provide an opportunity to observe inconsistent behavior. We exclude incomparable observations from our data in such cases.

Table 5 indicates that our subjects showed the same mixture of preferences between ambiguity and risk as those in the Hey et al. (2007) study when we exclude the pricing decision. As a robustness check, we note again that the risky gamble may have been a focal point for subjects who were indifferent and that some subjects reversed preferences. Neither of these factors affects our comparability with Hey et al. If we exclude the subjects who priced the two gambles identically and

Table 5: Comparisons of inferred preferences to other experimental studies.

	Prefer Risky	Prefer Ambiguous		
This study (ignoring price)	60.0%	40.0%		
Hey et al. (2007) (S1 model)	60.0%	40.0%		
	Prefer Risky	Prefer Ambiguous	No Preference	
This study (excluding inconsistent)	32.6%	19.6%	47.8%	
Fellner (1961)	33.3%	0.0%	66.7%	
Einhorn and Hogarth (1986)	47.0%	19.0%	34.0%	
Fox and Tversky (1995)(study 3)	60.4%	12.3%	27.4%	
	Prefer Risky	Prefer Ambiguous	No Preference	Inconsistent
This study	25.0%	15.0%	36.7%	23.3%
Becker and Brownson (1964)	47.1%	2.9%	14.7%	35.3%

those whose pricing decisions were inconsistent with their choices, we have $15/24 = 62.5\%$ strictly preferring the risky gamble. If we only drop the inconsistent choices, we have $29/46 \approx 63.0\%$ choosing the Risky Penny and pricing it at least as high as the Ambiguous Penny. So we confirm the proportions of risky and ambiguous choices of Hey et al.

The study by Fellner (1961) had approximately the same proportion of subjects strictly preferring the risky gamble to the ambiguous one as we found. However, none of his subjects strictly preferred the ambiguous gamble. He only had 12 subjects, however, so our numbers would only predict about two subjects in his population to have preferred ambiguity to risk. Fellner determined preferences by varying prizes, so his study is essentially a pricing task. In our study, $22/60 \approx 36.7\%$ of the subjects priced the risky gamble higher, which again is comparable to the results of Fellner. We had $16/60 \approx 26.7\%$ price the ambiguous coin higher, so by this measure, we would have expected around three of Fellner's subjects to have preferred the ambiguous lottery.

Einhorn and Hogarth (1986) closely match our proportion of subjects preferring the ambiguous lottery. However, they find more subjects preferring risky gambles and fewer indifferent subjects than we found. Their subjects were unpaid, and as we show in Table 6, behavior of unpaid subjects in ambiguity studies seems to diverge from that of paid subjects, to some degree increasing the value of risky choices over ambiguous ones. But some of this difference seems to be due to other factors: Fox and Tversky (1995) use paid subjects but, like Einhorn and Hogarth, find fewer subjects who are indifferent and more who strictly prefer the risky lottery. Reversals also do not seem to explain this difference. Becker and Brownson (1964) adjust for inconsistent choices. Like the other studies,

theirs finds that more subjects strictly prefer the risky lottery than we find, and that fewer are indifferent. And like all the studies except for Einhorn and Hogarth, Becker and Brownson find fewer subjects who strictly prefer ambiguity than we find.

Table 6: Comparisons of price solicited to other experimental studies.

	Average Scaled Price for		Payment Mechanism
	Risky Penny	Ambiguous Penny	
This study	0.48	0.46	BDM
Fox and Tversky (1995) (study 1)	0.24	0.15	Unpaid
Fox and Tversky (1995) (study 2)	0.49	0.43	38 step elicitation
Fox and Tversky (1995) (study 3)	0.56	0.45	BDM, for the pair
Arló-Costa and Helzner (2009)	0.42	0.13	Unpaid, 3 gamble task
Halevy (2007)	0.53	0.44	BDM, 4 gamble task

Overall, the subjects priced the Ambiguous Penny lower than the Risky Penny, as in many other studies (see Table 6). The difference between the prices is quite small, which matches other experimental research using some truth elicitation mechanism such as BDM. When subjects price hypothetical lotteries, as in the work of Arló-Costa and Helzner (2009), prices tend to diverge from our findings.

This result is most striking in the studies by Fox and Tversky (1995), where their hypothetical payoffs (scaled by the size of the prize) are much lower than our findings and the relative gap between risky and ambiguous lotteries is much higher. By contrast, their paid subjects give results resembling our own. One caveat is that the magnitudes of payoffs in their hypothetical studies are considerably larger than their actual payoffs (a \$100 hypothetical prize vs. a \$20 real prize). This is likely to account for at least part of the lower pricing magnitudes in their hypothetical lottery (for more on this, see Holt and Laury 2002).

Our findings are thus broadly consistent with other studies on ambiguity, but we should point out that our subjects price ambiguous choices higher than do the subjects in all the other studies. This difference is small with studies involving paid subjects. However, combining this with the results in Table 5 gives some suggestive evidence about ambiguity aversion. Previous studies may have overstated ambiguity aversion by confounding unknown distributions (ambiguity in our sense) with vagueness and comparative ignorance. We defer this issue to future work.

5 Conclusion

We interpret ambiguity as the absence of a uniquely defined probability. This makes ambiguity differ fundamentally from related concepts such as vagueness. A contract that leaves some details open to interpretation is vague but not necessarily ambiguous. Research and development of a novel technology requires decisions to be made under ambiguity, even if the procedure is spelled out in excruciating detail. If one accepts this interpretation, then it is entirely straightforward to draw ambiguous random numbers in laboratory experiments. Indeed, doing so is as easy as drawing random numbers from any well-behaved probability distribution. In the body of the article, we show how to do this. We close with remarks on why our procedure works.

To draw random numbers under risk, one needs a well-defined cumulative distribution—not necessarily one with a closed-form expression or one that is easily written down. But however it may be arrived at, the cdf must be uniquely defined at every point in its support. Boole’s insight was to consider the contrapositive of this idea. The key is to associate ambiguity with something like $0/0$. To generate ambiguity, all one needs is a random process with a cdf that is a partial function. If at any point, the cdf can take on a value like $0/0$, then the distribution of the process is ambiguous at that point.

The task of generating ambiguity, at least in terms of objective probabilities, is then exactly that of defining a partial function. Once that is accomplished, an experimenter can draw as many ambiguously distributed numbers as desired and can disclose anything about the procedure. There is no risk of the subjects learning the distribution as there is nothing to learn.

Our technique takes this idea to its logical extreme and constructs a process for which the cdf is ambiguous at *every* point. To do this, we use realizations of a Cauchy distributed random variable to construct a second-order Cauchy process. We justify this on both methodological and aesthetic grounds. The mean of a Cauchy distribution is $\infty - \infty$, which is just as indefinite as Boole’s $0/0$. By drawing from Cauchy distributions whose parameters depend on other Cauchy draws, we create a process in which the location and scale are ill-defined. So this is a convenient method for constructing an ill-defined cdf. The simulations in Section 3.2 demonstrate the divergence of the cdf. Aesthetically, the choice of a Cauchy process is appealing because it arises as a ratio

of two normals with zero mean. Because the normal distribution can be approximated with coin flips, it follows that any tool that can flip a coin can also generate ambiguity. We pursue this in Appendix A.1, where we replicate our simulations using only mechanical randomization devices.

Even if facing objective ambiguity, a decision maker might assign subjective probabilities to different possible outcomes. Our experiment in Section 4 addresses whether subjects who observe the results of our procedure act as if they face unique probabilities. We use a design based on the Ellsberg two-color urn and compare our results with those of other experiments on ambiguity. Our subjects saw 100 empirical histograms from our technique, each of which came from 3,000 draws.

After 300,000 observations, our subjects still react to our technique, for the most part, as other subjects have reacted to ambiguity in other studies. The proportions of subjects preferring risk to ambiguity and those who were indifferent are in line with other results in the field. We control for order of presentation, for whether the large prize is awarded for a coin landing heads or tails, and for the screen placement of the gambles on the left or right; none of these matter.

There is some suggestive evidence that our subjects are less ambiguity averse than subjects in other studies. Most of this difference arises when we compare our results to those with unpaid subjects. The remaining differences may reflect that our design separates vagueness, ignorance, and ambiguity and only tests the latter. Our technique gives future researchers a way to explore this further.

A Robustness Issues

A.1 Physical randomization

The simulations in Section 3.2 were run using computer pseudo-random number generators. This is solely a technical convenience, as we now demonstrate that we can draw ambiguously distributed numbers using mechanical randomization devices. This illustrates two important features of our technique: first, our technique does not depend on idealized mathematical properties or on computer software. Below, our randomization devices are essentially coin flips. Second, because our technique

is implementable with mechanical devices, it can be made as immune to manipulation as any other randomization technique is.

Generating approximately Cauchy distributed random variables is entirely straightforward. Laha (1959) provides a large number of bivariate distributions $G(X, Y)$ for which the ratio X/Y has a Cauchy distribution. In particular, for any two independent normally distributed random variables X, Y with $E[X] = E[Y] = 0$, X/Y is Cauchy distributed with location 0. Accordingly, if one can generate independent normally distributed random variables and divide, then one can obtain Cauchy distributed random variables.⁶ It follows that one can generate approximately Cauchy distributed random variables by flipping coins, summing, and dividing.

In practice, however, the number of coin flips required is likely to be unfeasibly large. By the central limit theorem, we can approximate a normally distributed random variable as the sum of, say, 30 coin flips, assigning a value of 1 to each coin that lands heads and 0 to each that lands tails. Subtracting 15 (rather than 14.99 to avoid the risk of dividing by zero) then gives an approximately normally distributed random variable with a mean close to zero, and then taking the ratio of two variables generated in this way gives a random variable that is approximately Cauchy.

This procedure requires us to flip 60 coins for each draw. So coin flips become tedious. We instead used a quantum random number generator, which generates a large number of quantum bits (qbits) by firing particles at a semi-transparent mirror, then attempting to detect the particles. Under quantum theory, the distribution of qbits cannot be uniquely defined, unless one is willing to consider probabilities that can be negative or above 1; for details, see Mückenheim et al. (1986), Suppes and Zanotti (1991), or Kronz (2007). Thus in some sense, the qbits are already ambiguous. Observationally, however, the realized qbits are at least close to a Bernoulli distribution with probability 1/2. For discussion, see Calude (2004).

Once one has a procedure for generating approximately Cauchy random variables, it is straightforward to generate a second-order Cauchy process, as described in Section 3.1. The first 60 coin flips (or qbits) provide a Cauchy distributed random variable with location $x_0 \approx 0$. Call this realization z_1 . To generate the next term in the process, draw another Cauchy random variable, and then add

⁶See Wakker and Klaassen (1995) on this point, and also on analysis of draws from a stationary Cauchy distribution.




Figure 6: Draws using physical randomization devices, 10 outcomes.

z_1 to the result. This yields a Cauchy distributed random number with location $x_0 + z_1$. For the n^{th} draw, $n \geq 3$, generate a new Cauchy distributed random variable, multiply by the absolute value of z_{n-2} , multiply by ϕ (which we set to 0.01 in these simulations), and add z_{n-1} . This matches our procedure with pseudo-random numbers.

After generating a sequence of 30 observations in this way, we mapped the results to the interval $\{1, \dots, 10\}$. We then repeated the procedure two more times. The histograms are shown in Figures 6 through 8.

In the first replication, the most common realized value was 3 (making up 30% of the draws), which did not show up at all in either the second or third replication. Among all three replications, the number 2 was never drawn; 8 was only drawn once (in the second replication); 6 was drawn twice (once each in the first and second replications). At the opposite extreme, 24.4% of the draws (22/90) were equal to 4, and 35.6% (32/90) were equal to 9. Overall, these two outcomes accounted for 60% of the realizations.

As a final test, we mapped these three sequences to $\{0, 1\}$. The frequency of a one (“success”) was 23.3% in the first replication and 30% in the third. Both of these were significantly different from $1/2$ (with 99.5% confidence for replication one and with 95% confidence for replication three). For replication two, the success frequency was 53.3%, which is indistinguishable from $1/2$. In sum,




Figure 7: Another set of draws using physical randomization devices.




Figure 8: A third set of physically generated draws.

the success probability in the second replication is different from that in the first and third, even though we followed an identical procedure.

A.2 Ambiguity with some probabilistic information

Our examples have illustrated the case of complete ambiguity, but one can also use our technique with some probabilistic information, as in the three-color Ellsberg experiment. Ellsberg's subjects were first asked to choose between a lottery giving a prize with probability $1/3$ and another giving the prize with ambiguous probability in $[0, 2/3]$. In a second task, the subjects chose between a lottery giving a prize with probability $2/3$ and another giving the prize with probability in $[1/3, 1]$. Thus the first ambiguous urn necessarily awards nothing with probability $1/3$, while the second necessarily awards a prize with probability $1/3$.

To replicate the ambiguity in the three-color Ellsberg experiments, one can proceed as follows: generate an urn with 60 balls. Give each ball one of two labels, such as {black, yellow}, using the method in Section 3.2. Add another 30 balls to the urn, each with a different label. This gives an urn with an ambiguous *composition*.

Alternatively, generate an urn with an ambiguously constructed *ratio* of black and yellow balls, keeping the fraction of red balls at $1/3$. To do this, draw a single number b from $\{0, \dots, 60\}$, representing the number of black balls in the urn. Set the number of yellow balls at $60 - b$.

The choice of technique can be illustrated directly to subjects, for example by showing them empirical histograms, as in our experiment. Thus, while the issue of randomizing over compositions or ratios remains, the experimenter need not be vague in what he or she discloses to the subjects.

B Instructions and Quiz

Instructions for the Experiment

This is a computerized experiment in the economics of decision-making. This experiment will last less than an hour.

You are guaranteed to receive \$7 for showing up on time. By following the instructions carefully, you may earn between \$10 and \$30 besides the participation fee. The actual amount of additional money that you may earn will depend on your choices and upon chance. We will pay you in cash after the experiment ends. You will need to sign and date a compensation receipt form before you receive your payment.

There are some rules you must follow:

1. Do not talk to others at any time during the experiment.
2. You will use your computer to select decisions during the experiment. Do not use your mouse or keyboard to play around with the software running on your computer. If you unintentionally or intentionally close the software program running on your computer, we will ask you to leave. If this happens, you will receive only your \$7 for showing up.
3. If you have any questions during the experiment, please raise your hand. An experimenter will come to your location and answer your questions.

You are free to withdraw from the experiment at any time, for any reason. If you choose to do so, please raise your hand. In this case, we will pay you the \$7 participation fee as you leave.

Details of the Experiment

We will use two computer pennies in this experiment. A computer penny is a penny whose flips are generated by a computer program. The pennies are called Penny C and Penny M.

To give you a feel for the properties of the two pennies, we built a collection of bar charts for each penny, where each chart shows you the results of 3,000 flips of that penny. The red bar denotes the number of times the penny landed Heads out of 3,000 flips and the blue bar denotes the number of times the penny landed Tails out of 3,000 flips. There are 100 such charts for each penny. The two collections of charts are shown at the end of these instructions. **Please look at them now.**

You will make three decisions over the course of the experiment, but only one choice will result in payment. First, you will decide which of two pennies to flip. Next, you will make two decisions on how much to sell each of the pennies for. The computer will randomly select one of the three decisions for payment.

The experimental software requires everyone to make his or her first decision before anyone can make his or her second decision. However, your payment is not affected by others' decisions.

Paired Choice: Which Penny Do You Want to Flip?

For the first choice of the experiment, you will be presented with the two computer pennies. Your task is to pick the penny you would wish to have flipped if you were rewarded according to the outcome of the flip. If the computer selects this decision, then we will pay you depending on whether the penny you chose lands on Heads or Tails.

At the beginning of the experiment, before you make any choice, the computer will tell you how much we will pay you if the penny you pick lands on Heads, and how much we will pay you if the penny you pick lands on Tails. The computer will either say that you get \$30 for Heads and \$10 for Tails, or that you get \$10 for Heads and \$30 for Tails.

Pricing Choices: How Much Would You Sell Each Penny For?

For each of the pennies, assume there is a potential buyer to whom you can sell the penny. You can sell the penny for any amount between \$10 and \$30. The buyer will determine his buying price by randomly picking a number between \$10 and \$30. If the buyer's randomly determined price is

higher than your selling price, you will receive the buying price and your payment will not depend on the outcome of the penny flip. If not, you will keep the penny, and you will receive \$30 or \$10, depending upon the penny flip.

For example, imagine you offer to sell one of the pennies to the potential buyer for \$15. Then your payment will depend upon whether the computer generates a higher or lower amount, and perhaps upon whether the penny lands on Heads or Tails. In the first two examples shown in the table below, the computer randomly generates a price greater than your selling price of \$15, so you are paid the amount drawn, and the flip does not matter (as you sold the penny). In the last two examples, the computer generated a price less than \$15, so you did not sell the penny. Your payment is then determined by the flip. For illustration, imagine that before you chose a price the computer had said you get \$10 for Heads and \$30 for Tails.

Buyer's Price	Your Price	Flip	Your Payment
\$27.1	\$15.00	Not Applicable	\$27.10
\$16.4	\$15.00	Not Applicable	\$16.40
\$12.8	\$15.00	Tails	\$30.00
\$14.2	\$15.00	Heads	\$10.00

We will ask you for a price for both pennies. Note that if you select a price of \$10, you will always sell the penny for the buyer's random price, and if you select a price of \$30, you will always keep the penny.

How Payment Is Determined

Note that we will not pay you for all your choices, but will randomly select one of your choices to play.

If the computer selects your decision on which penny to flip, then the computer will remind you which penny you chose, flip the penny, and pay you based on the flip.

If the computer selects one of the pricing decisions, then your payment will be based on a randomly generated amount and possibly the penny flip. If the computer generates a buyer's price higher than the price you chose for the selected penny, you will be paid the buyer's price. If the buyer's

price is lower than the price you chose for the selected penny, then the computer will flip the penny, and pay you based on the flip.

Quiz on Instructions

Below, please write down your answers to the following questions. In a few minutes, an experimenter will review the correct answers with you privately. Raise your hand when you're ready for review or have questions.

1. During the experiment, how many choices will you make?
 - a. 1
 - b. 2
 - c. 3
2. How many choices will be used to determine your payment?
 - a. 1
 - b. 2
 - c. 3
3. If your penny choice is randomly selected, how is your payment determined?
 - a. By flipping the penny
 - b. Comparing the random price to your selling price
 - c. If the random price is higher than your price, you will receive the random price, else you continue on and flip the penny.
4. If one of your pricing choices is randomly selected, how is your payment determined?
 - a. By flipping the penny
 - b. Comparing the random price to your selling price
 - c. If the random price is higher than your price, you will receive the random price, else you continue on and flip the penny.
5. Suppose one of your pricing choices is selected for payment. The randomly picked price is \$27.60. How will you be paid if your price is \$12.00?
 - a. Paid \$12.00
 - b. Paid \$27.60
 - c. By flip of heads and tails.
6. Suppose one of your pricing choices is selected for payment. The randomly picked price is \$27.60. How will you be paid if your price is \$28.00?
 - a. Paid \$28.00
 - b. Paid \$27.60
 - c. By flip of heads and tails.

7. Suppose one of your pricing choices is selected for payment. The randomly picked price is \$21.80. How will you be paid if your price is \$24.00?
- Paid \$24.00
 - Paid \$21.80
 - By flip of heads and tails.
8. Suppose one of your pricing choices is selected for payment. The randomly picked price is \$21.80. How will you be paid if your price is \$16.00?
- Paid \$16.00
 - Paid \$21.80
 - By flip of heads and tails.

References

- Althoen, S C, L King, K Schilling. 1993. How long is a game of snakes and ladders? *The Mathematical Gazette* **78**(478) 71–76.
- Arló-Costa, Horacio, Jeffrey Helzner. 2009. Iterated random selection as intermediate between risk and uncertainty. Manuscript, Carnegie Mellon University and Columbia University. In the electronic proceedings of the 6th International Symposium on Imprecise Probability: Theories and Applications.
- Barnett, V D. 1966. Order statistics estimators of the location of the Cauchy distribution. *Journal of the American Statistical Association* **61**(316) 1205–1218.
- Becker, Gordon M, Morris H DeGroot, Jacob Marschak. 1964. Measuring utility by a single-response sequential method. *Behavioral Science* **9**(2) 226–232.
- Becker, Selwyn W, Fred O Brownson. 1964. What price ambiguity? or the role of ambiguity in decision-making. *Journal of Political Economy* **72**(1) 62–73.
- Berlekamp, Elwyn R, John H Conway, Richard K Guy. 1982. *Winning Ways for Your Mathematical Plays*. Academic Press.
- Binmore, Kenneth. 1995. Debayesing game theory. Dag Prawitz, Brian Skyrms and Dag Westerstahl, ed(s)., *Studies in Logic and the Foundations of Game Theory: Proceedings of the Ninth International Congress of Logic, Methodology and the Philosophy of Science*. **13**4927–945
- Binmore, Kenneth. 2007a. Can knowledge be justified true belief? Working paper, University College, London.
- Binmore, Kenneth. 2007b. Making decisions in large worlds. Working paper, University College, London.
- Binmore, Kenneth G. 2009. *Rational Decisions*. Princeton University Press.

- Boole, George. 1854. *An Investigation of the Laws of Thought*. Dover (reprinted 1958).
- Calude, Cristian S. 2004. Algorithmic randomness, quantum physics, and incompleteness. M Margenstern, ed., *Proceedings of the Conference "Machines, Computations and Universality" MCU'2004*, no. 3354, *Lecture Notes in Computer Science*, Springer, 1–17.
- Chipman, John S. 1960. Stochastic choice and subjective probability. Dorothy Willner, ed., *Decisions, Values, and Groups*, vol. I. Pergamon Press, 70–95.
- Chow, Clare Chua, Rakesh K Sarin. 2002. Known, unknown, and unknowable uncertainties. *Theory and Decision* **52**(2) 127–138.
- Davidson, Donald, Patrick Suppes, Sidney Siegel. 1957. *Decision Making: An Experimental Approach*. Stanford University Press.
- DeGroot, Morris H, Prem K Goel. 1980. Estimation of the correlation coefficient from a broken random sample. *The Annals of Statistics* **8**(2) 264–278.
- Einhorn, Hillel J, Robin M Hogarth. 1986. Decision making under ambiguity. *Journal of Business* **59**(4, Part 2) S225–S250.
- Ellsberg, Daniel. 1961. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* **75**(4) 643–669.
- Fellner, William. 1961. Distortion of subjective probabilities as a reaction to uncertainty. *Quarterly Journal of Economics* **75**(4) 670–689.
- Fierens, Pablo Ignacio, Leonardo Chaves Rêgo, Terrence L Fine. 2009. A frequentist understanding of sets of measures. *Journal of Statistical Planning and Inference* **139**(6) 1879–1892.
- Fischbacher, Urs. 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* **10**(2) 171–178.
- Fishburn, Peter C. 1993. The axioms and algebra of ambiguity. *Theory and Decision* **34**(2) 119–137.
- Fox, Craig R, Amos Tversky. 1995. Ambiguity aversion and comparative ignorance. *Quarterly Journal of Economics* **110**(3) 585–603.
- Gilboa, Itzhak, David Schmeidler. 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* **18**(2) 141–153.
- Gilboa, Itzhak, David Schmeidler. 1993. Updating ambiguous beliefs. *Journal of Economic Theory* **59**(1) 33–49.
- Girón, Francisco Javier, Sixto Ríos. 1980. Quasi-Bayesian behaviour: A more realistic approach to decision making? *Trabajos de Estadística e Investigación Operativa* **31**(1) 17–38.
- Good, Irving John. 1983. *Good Thinking: The Foundations of Probability and Its Applications*. University of Minnesota Press.

- Halevy, Yoram. 2007. Ellsberg revisited: An experimental study. *Econometrica* **75**(2) 503–536.
- Hanany, Eran, Peter Klibanoff. 2007. Updating preferences with multiple priors. *Theoretical Economics* **2**(3) 261–298.
- Hayashi, Takashi, Ryoko Wada. 2006. An experimental analysis of attitude toward imprecise information. Working paper, Department of Economics, University of Texas at Austin and Department of Economics, Keiai University.
- Hayashi, Takashi, Ryoko Wada. 2009. Choice with imprecise information: An experimental approach. *Theory and Decision*, forthcoming.
- Hey, John Denis, Gianna Lotito, Anna Maffioletti. 2007. Choquet OK? Discussion paper 07/12, University of York.
- Holt, Charles A, Susan K Laury. 2002. Risk aversion and incentive effects. *American Economic Review* **92**(5) 1644–1655.
- Hurwicz, Leonid. 1951. Optimality criteria for decision making under ignorance. Discussion paper, Statistics 370, Cowles Commission.
- Jevons, William Stanley. 1874. *Principles of Science*. Macmillan.
- Kahneman, Daniel, Amos Tversky. 1973. On the psychology of prediction. *Psychological Review* **80**(4) 237–251.
- Keynes, John Maynard. 1921. *A Treatise on Probability*. Rough Draft (reprinted 2008).
- Klibanoff, Peter, Massimo Marinacci, Sujoy Mukerji. 2005. A smooth model of decision making under ambiguity. *Econometrica* **73**(6) 1849–1892.
- Knight, Frank H. 1921. *Risk, Uncertainty, and Profit*. Houghton Mifflin.
- Kolmogorov, Andrei Nikolaevich. 1983. *On Logical Foundations of Probability Theory*, vol. 1021, *Lecture Notes in Mathematics*. Springer, 1–5.
- Kronz, Fred. 2007. Non-monotonic probability theory and photon polarization. *Journal of Philosophical Logic* **36**(4) 449–472.
- Laha, R G. 1959. On a class of distribution functions where the quotient follows the Cauchy law. *Transactions of the American Mathematical Society* **93**(2) 205–215.
- Maccheroni, Fabio, Massimo Marinacci, Aldo Rustichini. 2006. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* **74**(6) 1447–1498.
- Milnor, John W. 1954. Games against nature. Robert M Thrall, Clyde H Coombs, Robert L Davis, eds., *Decision Processes*. Wiley, 49–59.
- Mückenheim, Wolfgang, G Ludwig, C Dewdney, P R Holland, A Kyprianidis, J P Vigier, N Cufaro Petroni, M S Bartlett, E T Jaynes. 1986. A review of extended probabilities. *Physics Reports* **133**(6) 337–401.

- Papamarcou, Adrian, Terrence L Fine. 1991. Stationary and almost sure divergence of time averages in interval-valued probability. *Journal of Theoretical Probability* **4**(2) 239–260.
- Pitman, E J G, E J Williams. 1967. Cauchy-distributed functions of Cauchy variates. *The Annals of Mathematical Statistics* **38**(3) 916–918.
- Savage, L James. 1972. *The Foundations of Statistics*, rev. ed. Dover.
- Scheinkman, José, Blake LeBaron. 1989. Nonlinear dynamics and stock returns. *Journal of Business* **62**(3) 311–337.
- Seidenfeld, Teddy, Larry Wasserman. 1993. Dilation for sets of probabilities. *The Annals of Statistics* **21**(3) 1139–1154.
- Stern, Hal, Thomas M Cover. 1989. Maximum entropy and the lottery. *Journal of the American Statistical Association* **84**(408) 980–985.
- Suppes, Patrick. 1974. The measurement of belief. *Journal of the Royal Statistical Society, Series B (Methodological)* **36**(2) 160–191.
- Suppes, Patrick, Mario Zanotti. 1991. Existence of hidden variables having only upper probabilities. *Foundations of Physics* **21**(12) 1479–1499.
- Sutton, John. 2006. Flexibility, profitability and survival in an (objective) model of Knightian uncertainty. Working paper, London School of Economics.
- Wakker, Peter P. 2008. Uncertainty. Steven N Durlauf, Lawrence E Blume, eds., *The New Palgrave Dictionary of Economics*. Macmillan, 428–439.
- Wakker, Peter P, Marc P Klaassen. 1995. Confidence intervals for cost/effectiveness ratios. *Health Economics* **4**(5) 373–381.
- Yates, J Frank, Lisa G Zukowski. 1976. Characterization of ambiguity in decision making. *Behavioral Science* **21**(1) 19–25.

Economic Science Institute Working Papers

2008

08-08 Stecher, J., Lunawat, R., Pronin, K. and Dickhaut, J. Decision Making and Trade without Probabilities.

08-07 Dickhaut, J., Lungu, O., Smith, V., Xin, B. and Rustichini, A. A Neuronal Mechanism of Choice.

08-06 Anctil, R., Dickhaut, J., Johnson, K., and Kanodia, C. Does Information Transparency Decrease Coordination Failure?

08-05 Tila, D. and Porter, D. Group Prediction in Information Markets With and Without Trading Information and Price Manipulation Incentives.

08-04 Caginalp, G., Hao, L., Porter, D. and Smith, V. Asset Market Reactions to News: An Experimental Study.

08-03 Thomas, C. and Wilson, B. Horizontal Product Differentiation in Auctions and Multilateral Negotiations.

08-02 Oprea, R., Wilson, B. and Zillante, A. War of Attrition: Evidence from a Laboratory Experiment on Market Exit.

08-01 Oprea, R., Porter, D., Hibbert, C., Hanson, R. and Tila, D. Can Manipulators Mislead Prediction Market Observers?