

The domain where I work is a crossroads connecting cognitive and social psychology, behavioral economics, management science, some aspects of applied mathematics and analytical philosophy. Some of the traffic there carries some interesting debates about the nature of uncertainty. Is it all one thing? For instance, can every version of uncertainty be reduced to some form of probability? Or are there different kinds? Was Keynes right to distinguish between “strength” and “weight” of evidence?

Why is this question important? Making decisions under uncertainty in a partially learnable world is one of the most important adaptive challenges facing humans and other species. If there are distinct (and consequential) kinds, then we may need distinct ways of coping with them. Each kind also may have its own special uses. Some kinds may be specific to an historical epoch or to a culture.

On the other hand, if there really is only one kind of uncertainty then a rational agent should have only one (consistent) method of dealing with it. Moreover, such an agent should not be distracted or influenced by any other “uncertainty” that is inconsequential.

A particular one-kind view dominated accounts of rationality in economics, psychology and related fields in the decades following World War II. The rational agent was a Bayesian probabilist, a decision-maker whose criterion for best option was the option that yielded the greatest expected return. This rational agent maximizes expected outcomes. Would you rather be given \$1,000,000 right now for sure, or offered a fair coin-toss in which Heads nets you \$2,400,000 and Tails leaves you nothing? The rational agent would go for the coin-toss because the expected return is $1/2$ times \$2,400,000, i.e., \$1,200,000, which exceeds \$1,000,000. A host of experimental evidence says that most of us would choose the \$1M instead.

One of the Bayesian rationalist’s earliest serious challenges came from a [1961](#) paper by Daniel Ellsberg (Yes, *that* [Daniel Ellsberg](#)). Ellsberg set up some ingenious experiments demonstrating that people are influenced by another kind of uncertainty besides probability: Ambiguity. Ambiguous probabilities are imprecise (e.g., a probability we can only say is somewhere between $1/3$ and $2/3$). Although a few precursors such as Frank Knight ([1921](#)) had explored some implications for decision makers when probabilities aren’t precisely known, Ellsberg brought this matter to the attention of behavioral economists and cognitive psychologists. His goal was to determine whether the distinction between ambiguous and precise probabilities has any “behavioral significance.”

Here’s an example of what he meant by that phrase. Suppose you have to choose between drawing one ball from one of two urns, A or B. Each urn contains 90 balls, with the offer of a prize of \$10 if the first ball drawn is either Red or Yellow. In Urn A, 30 of the balls are Red and the remaining 60 are either Black or Yellow but the number of Yellow and Black balls has been randomly selected by a computer program, and we don’t know what it is. Urn B contains 30 Red, 30 Yellow, and 30 Black balls. If you prefer Urn B, you’re manifesting an aversion to not knowing precisely how many Yellow balls there are in Urn A-- i.e., *ambiguity aversion*. In experiments of this kind, most people choose Urn B.

But now consider a choice between drawing a ball from one of two bags, each containing a thousand balls with a number on each ball, and a prize of \$1000 if your ball has the number 500 on it. In Bag A, the balls were numbered consecutively from 1 to 1000, so we know there is exactly one ball with the number 500 on it. Bag B, however, contains a thousand balls with numbers randomly chosen from 1 to 1000. If you prefer Bag B, then you may have reasoned that there is a chance of more than one ball in there with the number 500. If so, then you are demonstrating *ambiguity seeking*. Indeed, most people choose Bag B.

Now, Bayesian rational agents would not have a preference in either of the scenarios just described. They would find the expected return in Urns A and B to be identical, and likewise with Bags A and B. So, our responses to these scenarios indicate that we are reacting to the ambiguity of probabilities as if ambiguity is a consequential kind of uncertainty distinct from probability.

Ellsberg's most persuasive experiment (the one many writers describe in connection with "[Ellsberg's Paradox](#)") is as follows. Suppose we have an urn with 90 balls, of which 30 are Red and 60 are either Black or Yellow (but the proportions of each are unknown). If asked to choose between gambles A and B as shown in the upper part of the table below (i.e., betting on Red versus betting on Black), most people prefer A.

	30	60	
Urn	Red	Black	Yellow
A	\$100	\$0	\$0
B	\$0	\$100	\$0

	30	60	
Urn	Red	Black	Yellow
C	\$100	\$0	\$100
D	\$0	\$100	\$100

However, when asked to choose between gambles C and D, most people prefer D. People preferring A to B and D to C are violating a principle of rationality (often called the "Sure-Thing Principle") because the (C,D) pair simply adds \$100 for drawing a Yellow ball to the (A,B) pair. If rational agents prefer A to B, they should also prefer C to D. Ellsberg demonstrated that ambiguity influences decisions in a way that is incompatible with standard versions of rationality.

But is it irrational to be influenced by ambiguity? It isn't hard to find arguments for why a "reasonable" (if not rigorously rational) person would regard ambiguity as important. Greater ambiguity could imply greater variability in outcomes. Variability can have a downside. I tell students in my introductory statistics course to think about someone who can't swim considering whether to wade across a river. All they know is that the average depth of the river is 2 feet. If the river's depth doesn't vary much then all is well, but great variability could be fatal. Likewise, financial advisors routinely ask clients to consider how comfortable they are with "volatile" versus "conservative" investments. As many an investor has found out the hard way, high average returns in the long run are no guarantee against ruin in the short term.

So, a prudent agent could regard ambiguity as a tell-tale marker of high variability and decide accordingly. It all comes down to whether the agent restricts their measure of good decisions to expected return (i.e., the mean) or admits a second criterion (variability) into the mix. Several models of risk behavior in humans and other animals incorporate variability. A well-known animal foraging model, the energy budget rule, predicts that an animal will prefer foraging options with smaller variance when the expected energy intake

exceeds its requirements. However, it will prefer greater variance when the expected energy intake is below survival level because greater variability conveys a greater chance of obtaining the intake needed for survival.

There are some fascinating counter-arguments against all this, and counter-counter-arguments too. I'll revisit these in a future post.

I was intrigued by Ellsberg's paper and the literature following on from it, and I'd also read William Empson's classic [1930](#) book on seven types of ambiguity. It occurred to me that some forms of ambiguity are analogous to internal conflict. We speak of being "in two minds" when we're undecided about something, and we often simulate arguments back and forth in our own heads. So perhaps ambiguity isn't just a marker for variability. It could indicate conflict as well. I decided to run some experiments to see whether people respond to conflicting evidence as they do to ambiguity, but also pitting conflict against ambiguity.

Imagine that you're serving on a jury, and an important aspect of the prosecution's case turns on the color of the vehicle seen leaving the scene of a nighttime crime. Here are two possible scenarios involving eyewitness testimony on this matter:

- A) One eyewitness says the color was blue, and another says it was green.
- B) Both eyewitnesses say the color was either blue or green.

Which testimony would you rather receive? In which do the eyewitnesses sound more credible? If you're like most people, you'd choose (B) for both questions, despite the fact that from a purely informational standpoint there is no difference between them.

Evidence from this and other choice experiments in my [1999](#) paper suggests that in general, we are "conflict averse" in two senses:

1. Conflicting messages from two equally believable sources are dispreferred in general to two informatively equivalent, ambiguous, but agreeing messages from the same sources; and
2. Sharply conflicting sources are perceived as less credible than ambiguous agreeing sources.

The first effect goes further than simply deciding what we'd rather hear. It turns out that we will choose options for ourselves and others on the basis of *conflict aversion*. I found this in experiments asking people to make decisions about eyewitness testimony, risk messages, dieting programs, and hiring and firing employees. Nearly a decade later, Laure Cabantous (then at the Université Toulouse) confirmed my results in a [2007](#) paper demonstrating that insurers would charge a higher premium for insurance against mishaps whose risk information was conflictive than if the risk information was merely ambiguous.

Many of us have an intuition that if there's equal evidence for and against a proposition, then both kinds of evidence should be accessed by all knowledgeable, impartial experts in that domain. If laboratory A produces 10 studies supporting an hypothesis (H, say) and laboratory B produces 10 studies disconfirming H, something seems fishy. The collection of 20 studies will seem more trustworthy if laboratories A and B each produce 5 studies supporting H and 5 disconfirming it, even though there are still 10 studies for and 10 studies against H. We expect experts to agree with one another, even if their overall message is vague. Being confronted with sharply disagreeing experts sets off alarms for most of us because it suggests that the experts may not be impartial.

The finding that people attribute greater credibility to agreeing-ambiguous sources than conflicting-precise sources poses strategic difficulties for the sources themselves. Communicators embroiled in a controversy face a decision equivalent to a Prisoners' Dilemma, when considering communications strategies while knowing that other communicators hold views contrary to their own. If they "soften up" by conceding a point or two to the opposition they might win credibility points from their audience, but their opponents could play them for suckers by not conceding anything in return. On the other hand, if everyone concedes nothing then they all could lose credibility.

None of this is taken into account by the one-kind view of uncertainty. A probability "somewhere between 0 and 1" is treated identically to a probability of "exactly 1/2," and sharply disagreeing experts are regarded as the same as ambiguous agreeing experts. But you and I know differently.