

Late in my childhood I discovered that if I was vague enough about my opinions no-one could prove me wrong. It was a short step to realizing that if I was *totally vague* I'd *never* be wrong. What is the load-bearing capacity of the elevator in my building? If I say "anywhere from 0 tons to infinity" then I can't be wrong. What's the probability that Samantha Stosur will win a Grand Slam tennis title next year? If I say "anywhere from 0 to 1" then once again I can't be wrong.

There is at least one downside to vagueness, though: Indecision. Vagueness can be paralyzing. If the elevator's designers' estimate of its load-bearing capacity had been 0-to-infinity then no-one could decide whether any load should be lifted by it. If I can't be more precise than 0-to-1 about Stosur's chances of winning a Grand Slam title then I have no good reason to bet on or against that occurrence, no matter what odds the bookmakers offer me.

Nevertheless, in situations where we don't know much but are called upon to estimate something, some amount of vagueness seems not only natural but also appropriately cautious. You don't have to look hard to find vagueness; you can find it anywhere from politics to measurement. When you weigh yourself, your scale may only determine your weight within plus or minus half a pound. Even the most well-informed engineer's estimate of load-bearing capacity will have an interval around it. The amount of imprecision in these cases is determined by what we know about the limitations of our measuring instruments or our mathematical models of load-bearing structures.

How would a sensible decision maker work with vague estimates? Suppose our elevator specs say the load-bearing capacity is between 3800 and 4200 pounds. Our "worst-case" estimate is 3800. We would be prudent to use the elevator for loads weighing less than 3800 pounds and not for ones weighing more than that. But what if our goal was to overload the elevator? Then to be confident of overloading it, we'd want a load weighing more than 4200 pounds.

A similar prescription holds for betting on the basis of probabilities. We should use the lowest estimated probability of an event for considering bets that the event will happen and the highest estimated probability for bets that the event won't happen. If I think the probability that Samantha Stosur will win a Grand Slam title next year is somewhere from $1/2$ to $3/4$ then I should accept any odds longer than an even bet that she'll win one (i.e., I'll use $1/2$ as my betting probability on Stosur winning), and if I'm going to bet against it then I should accept any odds longer than 3 to 1 (i.e., I'll use $1 - 3/4 = 1/4$ as my betting probability against Stosur). For any odds offered to me corresponding to a probability inside my $1/2$ to $3/4$ interval, I don't have any good reason to accept or reject a bet for or against Stosur.

Are there decisions where vagueness serves no distinct function? Yes: Two-alternative forced-choice decisions. A two-alternative forced choice compels the decision maker to elect one threshold to determine which alternative is chosen. Either we put a particular load into the elevator or we do not, regardless of how imprecise the load-bearing estimate is. The fact that we use our "worst-case" estimate of 3800 pounds for our decisional threshold makes our decisional behavior no different from someone whose load-bearing estimate is *precisely* 3800 pounds.

It's only when we have a third "middle" option (such as not betting either way on Sam Stosur) that vagueness influences our decisional behavior distinctly from a precise estimation. Suppose Tom thinks the probability that Stosur will win a Grand Slam title is somewhere

between $1/2$ and $3/4$ whereas Harry thinks it's exactly $1/2$. For bets on Stosur, Harry won't accept anything shorter than even odds. For any odds shorter than that Harry will be willing to bet against Stosur. Offered any odds corresponding to a probability between $1/2$ and $3/4$, Tom could bet either way or even decide not to bet at all. Tom may choose the middle option (not betting at all) under those conditions, whereas Harry never does.

If this is how they always approach gambling, and if (let's say) on average the midpoints of Tom's probability intervals are as accurate as Harry's precise estimates, it would seem that in the long run they'll do about equally well. Harry might win more often than Tom but he'll lose more often too, because Tom refuses more bets.

But is Tom missing out on anything? According to philosopher Adam Elga, he is. In his [2010](#) paper, "Subjective Probabilities Should Be Sharp," Elga argues that perfectly rational agents always will have perfectly precise judged probabilities. He begins with his version of the "standard" rational agent, whose valuation is strictly monetary, and who judges the likelihood of events with precise probabilities. Then he offers two bets regarding the hypothesis (H) that it will rain tomorrow:

- *Bet A* If H is true, you lose \$10. Otherwise you win \$15.
- *Bet B* If H is true, you win \$15. Otherwise you lose \$10.

First he offers the bettor Bet A. Immediately after the bettor decides whether or not to accept Bet A, he offers Bet B.

Elga then declares that any perfectly rational bettor who is sequentially offered bets A and B with full knowledge in advance about the bets and no change of belief in H from one bet to the next will accept at least one of the bets. After all, accepting both of them nets the bettor \$5 for sure. If Harry's estimate of the probability of H, $P(H)$, is less than .6 then he accepts Bet A, whereas if it is greater than .6 Harry rejects Bet A. If Harry's $P(H)$ is low enough he may decide that accepting Bet A and rejecting Bet B has higher expected return than accepting both bets.

On the other hand, suppose Tom's assessment of $P(H)$ is somewhere between $1/4$ and $3/4$. Bet A is optional because .6 lies in this interval, and so is Bet B. According to the rules for imprecise probabilities, Tom could (optionally) reject both bets. But rejecting both would net him \$0, so he would miss out on \$5 for sure. According to Elga, this example suffices to show that vague probabilities do not lead to rational behavior.

But could Tom rationally refuse both bets? He could do so if not betting, to him, is worth at least \$5 regardless of whether it rains tomorrow or not. In short, Tom could be *rationally vague* if not betting has sufficient positive value (or if betting is sufficiently aversive) for him. So, for a rational agent, vagueness comes at a price. Moreover, the greater the vagueness, the greater value the middle option has to be in order for vagueness to earn its keep. An interval for $P(H)$ from $2/5$ to $3/5$ requires the no-bet option to have value equal to \$5, whereas an interval of $1/4$ to $3/4$ requires that option to be worth \$8.75.

Are there situations like this in the real world? A very common example is medical diagnostic testing (especially in preventative medicine), where the middle option is for the patient to undergo further examination or observation. Another example is courtroom trials. There are trials in which jurors may face more than the two options of conviction or acquittal,

such as mental illness or diminished responsibility. The Scottish system with its Not Proven alternative has been the subject of great controversy, due to the belief that it gives juries a let-out from returning Guilty verdicts. Not Proven is essentially a type of acquittal. The charges against the defendant are dismissed and she or he cannot be tried again for the same crime. About one-third of all jury acquittals in Scotland are the product of a Not Proven verdict.

Can an argument be made for Not Proven? A classic paper by Terry Connolly ([1987](#)) pointed out that a typical threshold probability of guilt associated with the phrase “beyond reasonable doubt” is in the [0.9, 1] range. For a logically consistent juror a threshold probability of 0.9 implies the difference between the subjective value of acquitting versus convicting the innocent is 9 times the difference in the value of convicting versus acquitting the guilty. Connolly demonstrated that the relative valuations of the four possible outcomes (convicting the guilty, acquitting the innocent, convicting the innocent and acquitting the guilty) that are compatible with such a high threshold probability are counterintuitive. Specifically, “. . . if one does [want to have a threshold of 0.9], one must be prepared to hold the acquittal of the guilty as highly desirable, at least in comparison to the other available outcomes” (pg. 111). He also showed that more intuitively reasonable valuations lead to unacceptably low threshold probabilities.

So there seems to be a conflict between the high threshold probability required by “beyond reasonable doubt” and the relative value we place on acquitting the guilty versus convicting them. In a [2006](#) paper I showed that incorporating a third middle option with a suitable mid-range threshold probability can resolve this quandary, permitting a rational juror to retain a high conviction threshold and still evaluate false acquittals as negatively as false convictions. In short, a middle option like “Not Proven” would seem to be just the ticket. The price paid for this solution is a more stringent standard of proof for outright acquittal than any probability of guilt shy of .9, because Not Proven requires a standard of proof between .9 and some lower probability.

But what do people do? Two of my Honours students, Sara Deady and Lavinia Gracik, decided to find out. Sara and Lavinia both have backgrounds in law and psychology. They designed and conducted experiments in which mock jurors deliberated on realistic trial cases. In one condition the Not Proven, Guilty and Acquitted options were available to the jurors, whereas in another condition only the conventional two were available.

Their findings, published in our [2007](#) paper, contradicted the view that the Not Proven option attracts jurors away from returning convictions. Instead, Not Proven more often supplanted outright acquittals. Judged probabilities of guilt from jurors returning Not Proven were mid-range, just as the “rational juror” model said they should be. They acquitted defendants only if they thought the probability of guilt was quite low.

Vagueness therefore earns its keep through the value of the “middle” option that it invokes. Usually that value is determined by a tradeoff between two desirable properties. In measurement, the tradeoff often is speed vs accuracy or expense vs precision. In decisions, it's decisiveness vs avoidance of undesirable errors. In betting or investment, cautiousness vs opportunities for high returns. And in negotiating agreements such as policies, it's consensus vs clarity.