Rational Choice

Subjective Probability

David Emmanuel Gray

Carnegie Mellon University in Qatar

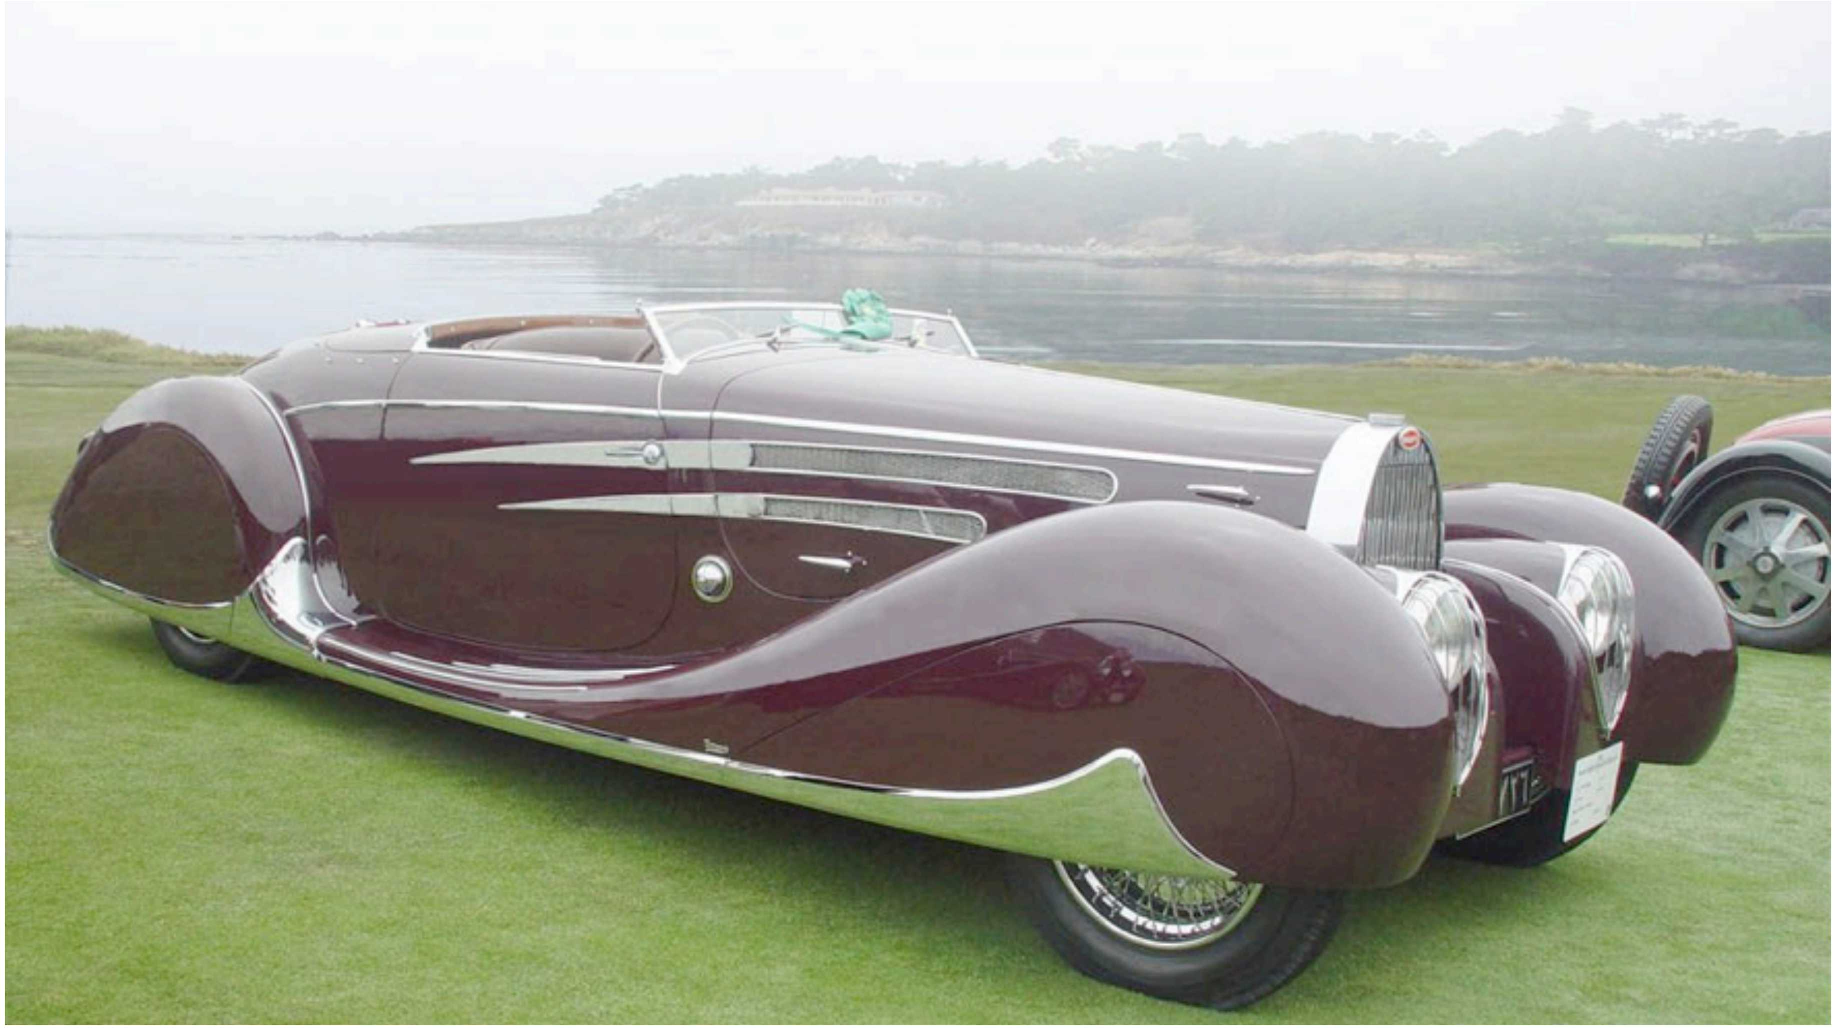

How likely do you think it is that your QR 20,000,000 Bugatti will be stolen this year?

Objective Probability

Objectivists would say this probability is a fact about the external world, though the exact way they would determine its value depends on their particular interpretation of probability:

1. The classical interpretation,
2. The frequency interpretation, or
3. The propensity interpretation.

We saw all of these last time.

Subjective Probability

Bayesians would say this probability is not about the external world. Your statement of the probability is just the degree to which you believe that your car might be stolen. The higher the number, the stronger your belief.

So long as these numbers obey the Kolmogorov axioms, your beliefs are rational. The Dutch book theorem (two classes from now) seeks to formally demonstrate this.

Subjective Probability

The key innovation of Bayesianism is to use your objectively observable behavior concerning risky prospects to model your probability distribution.

So your degree of belief concerning whether your car will be stolen this year can be ascertained by determining how much you are willing to pay to insure it against theft.

Buying Car Insurance

	Car is stolen	Car is not stolen
Buy insurance for QR x (Bet on the car being stolen)	Gain (value of car) Lose (value of car, insurance cost)	Gain (nothing) Lose (insurance cost)
Do not buy insurance (Bet against the car being stolen)	Gain (nothing) Lose (value of car)	Gain (nothing) Lose (nothing)

Buying Car Insurance

	Car is stolen	Car is not stolen
Buy insurance for QR x (Bet on the car being stolen)	–QR x	–QR x
Do not buy insurance (Bet against the car being stolen)	–QR 20,000,000	QR 0

Buying Car Insurance

Now the **fair** price of insurance (QR x) is the price at which you are indifferent between buying insurance and not buying it:

Buy insurance for QR x ~ do not buy insurance.

However, the price of insurance (QR x) is **favorable** when you prefer buying insurance to not buying it:

Buy insurance for QR x ≻ do not buy insurance.

Finally, the price of insurance (QR x) is **unfavorable** when you prefer not buying insurance to buying it:

Do not buy insurance ≻ buy insurance for QR x.

Buying Car Insurance

In order to determine your value for P(Car = Stolen), that is, the degree to which you believe that your car will be stolen this year, a Bayesian asks you to consider what you believe to be a *fair* price for the insurance.

Buying Car Insurance

Suppose you say QR 1,000,000 is a fair price. So:

Buy insurance for QR 1,000,000 ~ do not buy.

Assuming you maximize expected monetary value:

$$-\text{QR } 1{,}000{,}000 = P(\text{Car} = \text{Stolen}) \times (-\text{QR } 20{,}000{,}000) + (1 - P(\text{Car} = \text{Stolen})) \times (0).$$

Solving for P(Car = Stolen), P(Car = Stolen) = 0.05. You must therefore believe that there is a 5% chance that your car will be stolen this year.

Bayesian Calibration

Of course, this assumption rests on the fact that the utility for money is linear; whereas we have seen that the decreasing marginal utility for money may hold.

Consequently, Bayesians incorporate elements of von Neumann-Morgenstern utility theory in order to determine the utilities as well as the probabilities. The standard version of Bayesianism, devised by Leonard Savage, involves six axioms (see Box 7.1 in the text).

Bayesian Calibration

	E	$\bar{E}$
Bet x on E	gain (1 utile) lose (x utiles)	lose (x utiles)
Abstain	0 utiles	0 utiles

Bayesian Calibration

	E	$\bar{E}$
Bet x on E	$1 - x$	$-x$
Abstain	0	0

Bayesian Calibration

The bet on E is **fair** when the value of x is such that:

Bet x on E ~ abstain.

However, the bet on E is **favorable** when the value of x is such that:

Bet x on E ≻ abstain.

Finally, the bet on E is **unfavorable** when the value of x is such that:

Abstain ≻ bet x on E.

Bayesian Calibration

In order to determine your value for P(E), that is, the degree to which you believe that event E will occur, a Bayesian asks you to consider what value of x (in utiles) makes a bet on E *fair*.

Bayesian Calibration

Suppose you say $x = 0.3$ makes the bet on E fair. So:

Bet 0.3 on E ~ abstain.

Assuming you maximize expected utility:

$$0 = P(E) \times (1 - 0.3) + (1 - P(E)) \times (-0.3).$$

Solving for P(E), $P(E) = 0.3$. You must therefore believe that there is a 30% chance that E will occur.

We can do this for any event E, and $P(E) = x$.

Bayesianism and Conflict

A concern might linger about what it means for two Bayesians to disagree over the probability of a given event (i.e., they give different values for x in order to make a bet on E fair). Is at least one of these people wrong? Or are they both right?

Bayesians claim the question is irrelevant. They believe that "probability does not exist", and so it makes no sense to talk about right and wrong probabilities. All that matters is *coherence* (see the Dutch book theorem).

Bayesian Updating

Even so, if the two conflicting Bayesians use Bayes' theorem to update their beliefs (i.e., their probabilities), then these two individuals may, over time, converge on to the same probability distribution as they process more evidence.

Bayesian Updating

For instance, suppose you and I disagree over the probability that a coin biased in favor of heads, by coming up heads twice as likely as tails.

I believe there is a modest chance that the coin is biased in this way, so $P_{\text{David}}(\text{Coin} = \text{Biased}) = 0.30$. Meanwhile, you are pretty sure that the coin is *not* biased, so you say that $P_{\text{You}}(\text{Coin} = \text{Biased}) = 0.05$. So we have different **prior** probabilities concerning whether or not the coin is biased in this way.

Bayesian Updating

We *do* agree, however, on the following four conditional probabilities:

$$P(\text{Lands} = \text{Heads} \mid \text{Coin} = \text{Biased}) = \tfrac{2}{3},$$
$$P(\text{Lands} = \text{Tails} \mid \text{Coin} = \text{Biased}) = \tfrac{1}{3},$$
$$P(\text{Lands} = \text{Heads} \mid \text{Coin} \neq \text{Biased}) = \tfrac{1}{2}, \text{ and}$$
$$P(\text{Lands} = \text{Tails} \mid \text{Coin} \neq \text{Biased}) = \tfrac{1}{2}.$$

That is, we agree on what should happen *if* the coin is biased, and on what should happen *if* the coin is not biased. We just do not agree on the prior probabilities of those two big "ifs".

Bayesian Updating

The natural way to modify our beliefs is to test the coin. So suppose we flip the coin thirty times, and it comes up heads seven times and tails three times.

Since the coin flips are independent, we agree that

$$P(\text{Lands} = 7H3T \mid \text{Coin} = \text{Biased}) = (2/3)^7 \times (1/3)^3, \text{ and}$$

$$P(\text{Lands} = 7H3T \mid \text{Coin} \neq \text{Biased}) = (1/2)^7 \times (1/2)^3.$$

Bayesian Updating

Now we can use Bayes' theorem to update our *prior* probabilities into *posterior* probabilities:

$$P_{\text{David}}(\text{Coin} = \text{Biased} \mid \text{Lands} = 7\text{H}3\text{T})$$

$$= \frac{0.30 \times \left[(\tfrac{2}{3})^7 \times (\tfrac{1}{3})^3\right]}{\left\{0.30 \times \left[(\tfrac{2}{3})^7 \times (\tfrac{1}{3})^3\right]\right\} + \left\{(1 - 0.30) \times \left[(\tfrac{1}{2})^7 \times (\tfrac{1}{2})^3\right]\right\}} \approx 0.488, \text{ and}$$

$$P_{\text{You}}(\text{Coin} = \text{Biased} \mid \text{Lands} = 7\text{H}3\text{T})$$

$$= \frac{0.05 \times \left[(\tfrac{2}{3})^7 \times (\tfrac{1}{3})^3\right]}{\left\{0.05 \times \left[(\tfrac{2}{3})^7 \times (\tfrac{1}{3})^3\right]\right\} + \left\{(1 - 0.05) \times \left[(\tfrac{1}{2})^7 \times (\tfrac{1}{2})^3\right]\right\}} \approx 0.105.$$

Bayesian Updating

If we keep flipping the coin for many, many times, it is likely that we will converge onto a single posterior probability concerning whether the coin is biased. Ultimately, our respective original prior beliefs are "washed out". In this way, two disagreeing Bayesians can come to an agreed-upon, shared set of beliefs.

Next Class…

We will look at some problems with Bayesianism.