

What is Utility Measuring?

We're interested in what makes actions *choice-worthy* in the sense that's relevant to instrumental rationality. There are at least three broad accounts of what is in the best interests of an agent worth thinking about.

(I) **Hedonism**: What is most pleasurable, and least painful, is in an agent's greatest interest.

The main virtue of hedonism is welfare is made to be simple, roughly quantifiable, and easily intelligible. But it has troubles...

(II) **Objective-List Theory**: There is a list of goods (perhaps a different list for each agent), such that what makes an agent's life go well is that they achieve the things on that list.

This theory can explain why goods of life are heterogeneous. The theory is also incredibly flexible. But seems to be a single big obvious worry...

Finally, the favored view of decision theorists...

(III) **Preference-Satisfaction Theory**: An agent's life goes best to the extent that their strongest preferences are satisfied.

As with Hedonism, welfare is simple, roughly quantifiable, and intelligible. It avoids the problems with Hedonism because sometimes your preferences won't be fulfilled, even when you *think* they are. It also deals with “weighing” problems with ease: all weights are settled by preferences without further ado. As always, though, the theory faces problems of its own...

A Case Study in Locating Utilities: the “Paradox of Voting”

Sometimes it's tricky to pinpoint what people's preferences really are, and why they are what they are. If you should maximize expected utilities, there's a question about whether it's ever rational to aim at many things we do that have very low probabilities attached to them (as when buying a lottery ticket, e.g.).

Consider voting in the presidential election. Suppose the reason you vote is to get the person you vote for into office. How unlikely is it that your vote could decide the next presidential election? A study for the 2008 election estimated the chances of a single vote in New York State deciding the election at less than 1/1,000,000,000 (keep in mind: that's the *lower* bound, and the 2008 election was a *close* election.)

Question: How do we explain why people vote, given those odds? There seem to be three options (granting the rule of maximizing expected utility):

- (1) Voters *really* prefer their candidate over the other guy (this is especially surprising, given the median voter theorem, which we'll discuss later).
- (2) Voters aren't just voting to get their guy in. There's another outcome of their voting that they have (very?) strong preferences about.
- (3) Voters are confused.

(3) isn't entirely crazy. Let's make the this option vivid.

The odds of dying in a car accident in a given year in the US are about 1/65,000
(See the National Safety Council report for the bleak statistic.)

Let's guess the average number of car trips taken in year by someone in the US.
What I believe to be a conservative guess for our purposes: $365 \times 6 = 2190$

Putting these together we get an estimate of the likelihood of dying on any car trip:

$$1/142,000,000$$

That means a driver in upstate New York seems to be (at least) seven times as likely to die on her way to the polling station as she is to change who wins the election by voting. Do you value getting someone into office seven times as much as *not dying*?

Can we say something good about (1) or (2)?