

TWO SELVES

The term *utility* has had two distinct meanings in its long history. Jeremy Bentham opened his *Introduction to the Principles of Morals and Legislation* with the famous sentence “Nature has placed mankind under the governance of two sovereign masters, *pain* and *pleasure*. It is for them alone to point out what we ought to do, as well as to determine what we shall do.” In an awkward footnote, Bentham apologized for applying the word *utility* to these experiences, saying that he had been unable to find a better word. To distinguish Bentham’s interpretation of the term, I will call it *experienced utility*.

For the last 100 years, economists have used the same word to mean something else. As economists and decision theorists apply the term, it means “wantability”—and I have called it *decision utility*. Expected utility theory, for example, is entirely about the rules of rationality that should govern decision utilities; it has nothing at all to say about hedonic experiences. Of course, the two concepts of utility will coincide if people want what they will enjoy, and enjoy what they chose for themselves—and this assumption of coincidence is implicit in the general idea that economic agents are rational. Rational agents are expected to know their tastes, both present and future, and they are supposed to make good decisions that will maximize these interests.

EXPERIENCED UTILITY

My fascination with the possible discrepancies between experienced utility and decision utility goes back a long way. While Amos and I were still working on prospect theory, I formulated a puzzle, which

went like this: imagine an individual who receives one painful injection every day. There is no adaptation; the pain is the same day to day. Will people attach the same value to reducing the number of planned injections from 20 to 18 as from 6 to 4? Is there any justification for a distinction?

I did not collect data, because the outcome was evident. You can verify for yourself that you would pay more to reduce the number of injections by a third (from 6 to 4) than by one tenth (from 20 to 18). The decision utility of avoiding two injections is higher in the first case than in the second, and everyone will pay more for the first reduction than for the second. But this difference is absurd. If the pain does not change from day to day, what could justify assigning different utilities to a reduction of the total amount of pain by two injections, depending on the number of previous injections? In the terms we would use today, the puzzle introduced the idea that experienced utility could be measured by the number of injections. It also suggested that, at least in some cases, experienced utility is the criterion by which a decision should be assessed. A decision maker who pays different amounts to achieve the same gain of experienced utility (or be spared the same loss) is making a mistake. You may find this observation obvious, but in decision theory the only basis for judging that a decision is wrong is inconsistency with other preferences. Amos and I discussed the problem but we did not pursue it. Many years later, I returned to it.

EXPERIENCE AND MEMORY

How can experienced utility be measured? How should we answer questions such as “How much pain did Helen suffer during the medical procedure?” or “How much enjoyment did she get from her 20 minutes on the beach?” The British economist Francis Edgeworth speculated about this topic in the nineteenth century and proposed the idea of a “hedonimeter,” an imaginary instrument analogous to the devices used in weather-recording stations, which would

measure the level of pleasure or pain that an individual experiences at any moment.

Experienced utility would vary, much as daily temperature or barometric pressure do, and the results would be plotted as a function of time. The answer to the question of how much pain or pleasure Helen experienced during her medical procedure or vacation would be the “area under the curve.” Time plays a critical role in Edgeworth’s conception. If Helen stays on the beach for 40 minutes instead of 20, and her enjoyment remains as intense, then the total experienced utility of that episode doubles, just as doubling the number of injections makes a course of injections twice as bad. This was Edgeworth’s theory, and we now have a precise understanding of the conditions under which his theory holds.

The graphs in [figure 15](#) show profiles of the experiences of two patients undergoing a painful colonoscopy, drawn from a study that Don Redelmeier and I designed together. Redelmeier, a physician and researcher at the University of Toronto, carried it out in the early 1990s. This procedure is now routinely administered with an anesthetic as well as an amnesic drug, but these drugs were not as widespread when our data were collected. The patients were prompted every 60 seconds to indicate the level of pain they experienced at the moment. The data shown are on a scale where zero is “no pain at all” and 10 is “intolerable pain.” As you can see, the experience of each patient varied considerably during the procedure, which lasted 8 minutes for patient A and 24 minutes for patient B (the last reading of zero pain was recorded after the end of the procedure). A total of 154 patients participated in the experiment; the shortest procedure lasted 4 minutes, the longest 69 minutes.

Next, consider an easy question: Assuming that the two patients used the scale of pain similarly, which patient suffered more? No contest. There is general agreement that patient B had the worse time. Patient B spent at least as much time as patient A at any level of pain, and the “area under the curve” is clearly larger for B than for A. The key factor, of course, is that B’s procedure lasted much

longer. I will call the measures based on reports of momentary pain hedonimeter totals.




Figure 15

When the procedure was over, all participants were asked to rate “the total amount of pain” they had experienced during the procedure. The wording was intended to encourage them to think of the integral of the pain they had reported, reproducing the hedonimeter totals. Surprisingly, the patients did nothing of the kind. The statistical analysis revealed two findings, which illustrate a pattern we have observed in other experiments:

- **Peak-end rule:** The global retrospective rating was well predicted by the average of the level of pain reported at the worst moment of the experience and at its end.
- **Duration neglect:** The duration of the procedure had no effect whatsoever on the ratings of total pain.

You can now apply these rules to the profiles of patients A and B. The worst rating (8 on the 10-point scale) was the same for both patients, but the last rating before the end of the procedure was 7 for patient A and only 1 for patient B. The peak-end average was therefore 7.5 for patient A and only 4.5 for patient B. As expected, patient A retained a much worse memory of the episode than