Utility

Utility, a concept invented by Jeremy Bentham to mean something close to pleasure or happiness. According to his philosophy of utilitarianism one ought to act to maximize the total, or perhaps the average, of human utility, somewhat misleadingly described as "the greatest good for the greatest number."

In economics, however, your utility function describes not what you enjoy doing but what you choose to do. "The utility to me of consuming an apple is greater than that of consuming an orange" means that, given the choice, I will choose the former over the latter. We usually expect people to choose what makes them happy, hence expect a close correlation between utility in the economist's sense and utility in the philosopher's sense. That matters because one of the things economists do, when they are not making a point of being objective, value-free scientists, is to draw conclusions about what people ought to do, for instance that they ought to abolish tariffs and price controls.[1] Those conclusions usually depend on the assumption, stated or unstated, that maximizing utility in the economist's sense will also maximize it in the philosopher's sense. That was clearer a little over a century ago when the economic arguments were being made by an economist, Alfred Marshall, who was not afraid to make explicit the utilitarian foundations of his economic conclusions.

The two senses of utility are correlated but not the same. Suppose you are going to die six months from now. Is your utility greater if you have several months advanced warning, as cancer patients often do, or if your death comes as a complete surprise? Spending several months knowing that you are about to die would be, for most of us, a very unpleasant experience. If utility is another word for happiness, imagined as a characteristic of what is going on inside your head, the second alternative is almost certainly preferable to the first.

Happiness is not all that matters to people. If one could choose in advance, perhaps by instructions to your doctor, many of us would choose to know. We have things we would like to get done before dying, things to be said to children, wife, friends, projects to be completed whose completion matters, if only to our sense of having lived a life worth living, arrangements to be made for the future of those dear to us. A close friend, not all that long ago, spent a good deal of his last few months reducing to something more like order his crowded and cluttered house for the benefit of his wife and daughters.

For another example of tension between what makes you happy and what you choose, consider Robert Nozick's experience machine.[2]

> Someone invents an experience machine; get into it and you will have a fully convincing illusion of experience. The inventor, who somehow knows what your life is going to be like, makes you the following offer:
>
> *Get into my experience machine, spend the rest of your life there, and I will give you the illusion of a life slightly better than the one you would otherwise live. Your average income in the illusion will be a few thousand dollars higher than it would have been in reality, your wife a little prettier, your children slightly better behaved, your promotions*

---

[1] I discuss the issue in Chapter 15 of *Price Theory*.

[2] Nozick describes the experience machine in *Anarchy, State and Utopia*. The version of the idea presented here is based on his but modified for my purposes.

> *just a little prompter. Your illusory summers won't be quite as hot, or winters quite as cold. Once you are in the machine you will not know that it is an illusion.*

Assume you believe his offer. Do you accept it? If not, why not?

The answer depends on what you value. If all that all that matters to you about the world is how it impinges on you, what effect it has on your sensations, you should accept the offer. But for me and, I suspect, many other people, that is not all that matters. I do not merely want the illusion of having written an interesting, enjoyable and original book, I want to have actually done it. I do not just want to think people read my books and are affected by the ideas in them, I want them to actually read and be affected. I don't want just the illusion of wonderful children, I want my wonderful children to actually exist.

I wouldn't touch that machine with a ten-foot pole.

The question applies to things more realistic than Nozick's hypothetical. Consider recreational drugs. A lot of us have a gut level feeling that the pleasure from being high on a drug is somehow less valid, less real, than the pleasure from accomplishing something, even if it is only winning a game of tennis or climbing a mountain. Feeling good about yourself because you are drunk is somehow less valid than feeling good about yourself because you have just saved someone's life at risk of your own or solved an important problem.

Consider virtual reality, which gets us closer to Nozick's experience machine. Why do I feel better about making Germanic lyres in my basement than about making Whitesoul Helms in *World of Warcraft*? Why do I feel less comfortable about consuming many hours online fighting computer-generated monsters than about consuming a similar amount of time, also online, arguing with people about subjects of interest to me?

The problem is older than *World of Warcraft*. I know some very smart people who put substantial amounts of time and effort into playing games: chess, bridge, poker. As long as one views it as recreation, there is no problem. But what about a really able person who treats the game as his real life and whatever he does to earn money for food and rent as an annoying distraction? Something about that feels wrong, feels as though he would be leading a better life if he put the same talent and passion into building better houses or writing better computer programs.

Or perhaps not.

Other Worlds and Wasted Talents

Over the course of my life, I have spent a good deal of time in what I think of as other worlds — The Society for Creative Anachronism, *World of Warcraft*, interacting online in Usenet news groups, and the like. One thing that strikes me about such worlds is the presence of competent, energetic, talented people whose life in the ordinary world reflects little or none of that.

I was reminded of this when someone I knew in *WoW* as an unusually competent and charismatic leader, organizer, and player, mentioned the problem of "parental agro." He was apparently a college student, possibly a graduate student, living with his parents. Older examples are friends in the SCA of whose abilities and energy I think highly, who made their living as school teachers or secretaries or the like — respectable jobs, but not ones that make much use of those abilities.

It is not surprising that an energetic individual who does not find much outlet for his energies in his current career will direct them towards his hobbies. Adam Smith long ago observed that, in the

British universities of the time, a professor got no benefit by doing a good job of teaching, since the professors were on salary rather than, as in at least some of the Scottish universities, paid by his students. He concluded that if the professor were naturally energetic he would spend his energies doing something that might be of some benefit to him rather than his job, which would not. Nowadays we call it "consulting."

And yet, putting all of your talent into play seems wasteful. Starting a business, running a restaurant, doing scientific research, have the same potential for employing human talents as organizing a guild in WoW or an event in the SCA. They also produce other benefits, most obviously the opportunity to combine fun and profit in a way rarely possible with one's hobbies. I am glad that these people spend the time and talents they do in worlds we share, since I benefit from their doing so. But I wonder what keeps them from employing the same talents more successfully in the part of the world where they spend forty hours a week making a living.

Is it that they prefer for that part of their life to make fewer demands on them? Or is it rather a case of wasted talents, a failure of current institutions to do as good a job as they might of letting potentially productive individuals find suitable employment for their abilities?

Perhaps I am looking at it backwards. It seems odd that someone with the talent to be successful in play would not be successful in work. It is not at all odd that someone who, for whatever reason, is unable to find a rewarding career chooses to get his rewards elsewhere.

Should Irrational Preferences Count?

The field of behavioral economics discussed in Chapter XXX[Beyond Rationality] deals with predictable patterns of behavior that appear inconsistent with rationality as economists understand it. My one contribution to the field is a chapter, "Economics and Evolutionary Psychology," in the book *Evolutionary Psychology and Economic Theory*; a draft is available on my web page. In it I try to show that several patterns of behavior which are puzzling in terms of the assumptions of economics make sense in terms of evolutionary psychology; they can be explained as behavior that got hardwired into us because it increased an individual's reproductive success in the hunter gatherer societies where our species spent most of its history.

Consider, as one example, the endowment effect, the observation that individuals value items that belong to them more than items that do not, even if, as in the classic Cornell coffee cup experiment, who owns what is the result of random chance. I explain this as a commitment strategy designed to enforce property rights in a world without police and courts, the human elaboration of the territorial behavior observed in many animal species. If I am willing to fight harder to defend what is mine than to take what is yours, you are similarly willing, and both of us know it, then most of the time ownership will not have to be fought for.

Economists define efficiency in terms of people getting what they want, not what we think they should want. Suppose you accept my explanation for the endowment effect or a similar explanation for some other apparently irrational pattern of behavior, that it exists not because it serves the present interest of the individual but because it served the reproductive interest of other individuals long ago in a very different environment. Should you still take that preference as a given in evaluating economic institutions?

Before answering, you might want to consider a simpler question of the same sort. You observe A add some cyanide to B's wine glass while B is looking in the other direction. You ask B if he wants to drink what is in the glass. B replies that he does. Do you conclude that his drinking it is, on the principle of revealed preference, a good thing?

A possible response is that what B wants is wine, not cyanide; if he knew there was cyanide in the wine glass he would not want to drink it. I am not sure that solves the problem. After reading my article and being convinced by it of the source of the endowment effect, you might still value something more because you have it.

I suspect I do.

Why Risk Aversion Isn't

Many fields use technical terms that sound self-explanatory and aren't. The result is that many people believe they know what those terms mean and don't. I am confident that there are millions of people who believe that they understand what the Theory of Relativity says, even if not the mathematical details. The theory says that everything is relative. Surely that is clear enough.

Clear but wrong. One of the many things the Theory of Relativity tells us is that the speed of light is quite impossibly absolute, the same relative to you however fast you are moving. Economics has similar problems with terms such as *efficiency* and *competition*.

Another example is *risk aversion*. It sounds as though it means aversion to risk, so one might expect a risk averse person to avoid dangerous hobbies, a risk preferring person to be drawn to them. It is not true; there is nothing in the definition of risk aversion that implies that a risk averse person is less likely to take up hang gliding or mountain climbing than a risk preferer.

The definition of risk aversion, as any good textbook that covers the subject will explain, is that a risk averse person, faced with the choice between an uncertain set of monetary payments and a certain payment with the same expected value, will prefer the latter. As that definition suggests, it is a statement not about his taste for risk but about his taste for money.

To see why we would expect people to be risk averse, imagine that you are faced with two possible jobs. One pays you $60,000/year. The other has equal odds of paying you $20,000/year or $100,000 year. We expect most people to prefer the former job, all else being equal. Why?

If you accept the risky alternative, you are giving up dollars in the future where you lose the bet in order to get dollars in the future where you win the bet. You are giving up (probabilistic) dollars used to buy things you would get as your income increased from $20,000 to $60,000 in order to get (probabilistic) dollars to buy things you would get as it increased from $60,000 to $100,000. As your income increases you buy the more important things first, so we would expect the gain from getting a dollar at the high end to be less than the loss from losing one at the low.

As this (entirely conventional) exposition shows, risk aversion is simply declining marginal utility of income. The fact that your marginal utility of income decreases as your income increases tells us nothing at all about how the marginal utility of other things changes as the amount you have of them changes, hence the fact that you are risk averse does not tell us what your attitude will be to risks that involve non-monetary payoffs.

Your doctor calls you into his office to give you some very bad news. You have been diagnosed with a disease that, if untreated, will kill you in fifteen years. There is an operation which, if it succeeds, will let you live thirty years, but half the time instead kills the patient. You have a choice of a certainty of fifteen years or a fifty/fifty gamble between thirty and zero.

As it happens, the one thing in life you most want to do is to produce and bring up children. Thirty years is long enough to do that; fifteen is not. You grit your teeth and sign up for the operation. You are risk preferring in years of life, because years of life have increasing marginal utility to you. You may be, probably are, risk averse in dollars because dollars have decreasing marginal utility to you.

Risk preference, as economists use the term, is not about risk.

Declining marginal utility, progressive taxes, and fairness

I mentioned in another chapter an exchange I had with Scott Alexander, a blogger I think highly of, on the subject of taxation. He agreed that, on average, richer people in our system pay a larger fraction of their income as taxes, but argued that fairness required even greater progressivity. Asked to explain and defend the basis for that belief, he offered the usual argument for declining marginal utility of income, claimed that fairness required equal utility burdens on rich and poor, and concluded that the tax system ought to be highly graduated.

His argument, if true, implies that richer people should pay more dollars in taxes than poorer people. It does not tell us whether they should pay a larger or smaller proportion of their income. Consider two taxpayers, one with an income of $40,000/year, one with an income of $80,000. Consistent with declining marginal utility of income, assume that the former has a marginal utility of income of two utiles/dollar, the latter of one utile/dollar. Assume a flat tax which collects $4,000 from the poorer taxpayer, $8,000 from the richer. The utility cost of the tax is then 8000 utiles for each, hence "fair" by the standard of equal utility burden. The utility cost to the richer person of each dollar he pays is half as much but he is paying twice as many dollars.[3]

Generalizing this example, we can see that if marginal utility of income declines with increasing income less than proportionally, — if, say, MU(I)=AI^(-.9) — then equal utility shares imply a regressive tax, while if it declines more than proportionally to income, the same rule implies a progressive tax.

The next interesting question is whether the rule itself makes any sense. I do not see that it does. I can see a philosophical argument for the claim that everyone should end up with the same income, although I am not convinced by it. I can see a utilitarian argument for redistribution designed to transfer income from those with low marginal utility of income to those with high marginal utility of income, although I can also see utilitarian problems with such a policy. But a rule of equal utility cost faces two obvious problems. The first is that taxpayers do not all get the same utility benefit from the state, the second that the state does not get the same benefit from all taxpayers.

Consider a program such as social security which collects money and pays out money. Dollars collected from the richer taxpayer probably cost him less utility than dollars collected from the poorer taxpayer cost him. But dollars paid to the richer taxpayers also provide less utility than dollars paid to the poorer. So a rule of equal utility burden means that the poor are getting, in utility terms, a much better deal than the rich.

---

[3] Mathematically sophisticated readers may object that I am only specifying the marginal utility of the last dollar, but should be able to show that if the pattern holds throughout the income range, with MU(I) = A/I, the conclusion still holds.

I do not think there is such a thing as a fair tax, even if we are willing to separate questions of fairness from questions of justice. But many people do believe that there is such a thing and that they know about what it would be.

It's a harder project than they may suppose.