Behavioral Economics

Mark Dean

Lecture Notes for Fall 2009 Introductory Microeconomics - Brown University

1 Introduction

We will finish the course with a couple of observations about behavioral economics. Behavioral economics is a (relatively) recent development in economics which uses insights from experimental economics and psychology to critique the standard models of economics behavior that we have discussed so far. The idea is that, by improving our models of economic behavior, we will improve our understanding of the way that economies work.

One word of warning before we begin - in order to be useful, it is not enough for the behavioral economist to show that existing models are wrong. Of **course** they are wrong. As we discussed at the start of the course, the world is a complicated place, and human behavior is particularly complicated. We can only hope that our models are crude, but useful approximations of the way that people actually behave. In my opinion, a behavioral observation has to fulfil (at least) the following two criteria to be interesting

1. It must by a **systematic** deviation from the rational paridigm
2. This deviation must be important for some economic phenomena

Keep these criteria in mind as we discuss the behavioral anomalies below

Here, we are going to somewhat crudely split the phenomena that we are going to discuss into two distinct groups

1. Those that violate basic rationality - in other words, they do not act as if they are maximizing a fixed utility function

2. Those that satisfy basic rationality, but violate some other assumptions that are usually made in economic modelling

2 Violations of Basic Rationality

A fundamental tenant of standard economic theory is that people make choices in order to maximize a stable utility function. As we discussed in class, this is basically the same as assuming that people satisfy the independence of irrelevant alternatives. This assumption is important not only for our predictions of the way in which people behave, but also for our notion of welfare. All the standardly used measures of welfare in economics are based on preferences that are revealed by the choices that people make. If these preferences are not consistent, then such measures are meaningless.

How well does the idea of fixed, stable preferences hold up in practice? Many of the critiques of this assumption fall under the heading of **framing effects:** seemingly inconsequential changes in the way that an option is presented can lead to big changes in the choices that people make. Perhaps the most famous example comes from Kahneman and Tversky, who compared people's answers to the following question:

Example 1 *imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume the exact scientific estimate of the consequences of the programs are as follow:*

Two different groups of people were then offered different alternatives. The first group was offered:

1. Program A: Exactly 200 people will be saved

2. Program B: There is a 1/3 probability that 600 people will be saved and a 2/3 probability that no-one will be saved

While the second group were offered

1. Program C: Exactly 400 people will die

2. Program D: There is a 1/3 probability that no-one will die and a 2/3 probability that 600 people will die

In most experiments people overwhelmingly choose A from the first group and D from the second group. This is taken as evidence that the same person would choose A from A and B and D from C and D. Why is this a problem? Well, if you re-read the choices carefully, you will see that the outcome in A is same as the outcome in C and the outcome in B is exactly the same as the outcome in D. Obviously not exactly the same, as the outcome are described, or framed differently, but in all important respects, the outcomes implied by the choices seem to be the same. This is an important, and worrying observation for economics, because it suggests that there is no 'revealed preference' between these two outcomes - people's choices do not appear to represent a stable underlying preference ordering.

What is driving this result? Well, Amos Tverskey and Daniel Kahneman (who proposed the above experiment) suggested that the important difference in the framing between the two problems was in what it was that subjects perceived as a loss. In the first case, the wording of the question sets the 'reference point' to be that everyone will die (i.e. the language is about people being 'saved'). Thus, any amount of lives saved seems like a 'gain' from that point of view. In contrast, in the second case, the reference point is that everyone will live (i.e. the language is about people dieing). Kahneman and Tverskey's idea was that people are *loss averse* - meaning that they really hate losing something for sure. In fact, so much do they hate it, that they become risk loving in choices between losses - just in the above case. The idea of losing 400 lives for sure is so horrible that people are prepared to take a gamble to avoid it. In contrast, in the 'gain domain', people are risk averse, as they are in standard theory. Thus the utility function is kinked at the reference point (as we will show in class).

Kahneman and and Tverskey formalized this idea in a behavioral model called 'prospect theory', which has subequently gone on to be one of the two most cited economics papers of all time, winning them the Nobel prize along the way. Despite this, and the fact that prospect theory is widely used, it remains a little controversial. Some behavioralists claim that people violate the tenants of prospect theory as often as standard expected utility theory. Note also that, while prospect theory can help economists to explain what people may do, it does not recover the idea of revealed preference -

peoples's choices are still inconsistent.

2.1 Status Quo Bias

Another important type of framing effect is status quo bias. This is the propensity for people to choose a particular alternative just because it is the status quo, or default option. Examples of status quo bias include

1. People demanding a higher price to sell an object that they own than they are prepared to pay for the same object
2. Higher enrollment rates in companies that have opt out 401k investment plans than those that have opt in plans
3. People being more likely to use internet explorer on machines where it is already installed

One of the potentially confusing things about current research into status quo bias is that the concept of status quo is often poorly defined. For the purposes of this note, I am going to use a broad definition of status quo, using it to refer to an alternative that has at least one of the following two properties:

1. It is the alternative that the decision maker is currently choosing
2. It is the alternative that the decision maker will end up with if they do not actively choose something else

Of course, in many cases, an alternative may have both of these properties.. However, it is not the case that the two always come together. When facing a new environment (such as signing up for a 401k plan for the first time) there may be a 'default' option that a decision maker will get if they do not actively choose something else (for example a default investment fund). Such an option would have property 2, but not property 1. Furthermore, one can think of cases in which the option that the decision maker is currently choosing is not what they will get if they do nothing – one example of this is the random reassignment of customers to different long distance carriers after the breakup of AT&T in 1984 [Schweitzer, 1994]. Thus an option can have property 1 but

not property 2. This latter scenario is arguably closely related to the remedy suggested in Case T-201/04, as I discuss below.

While many studies confound these two effects, there is at least one [Schweitzer, 1994] that explicitly attempts to find independent evidence for the each of them. This study, which uses laboratory experiments, finds evidence for status quo bias caused both by previous ownership (property 1) and by lack of need for an active choice (property 2). The study also finds that the two effects are likely to reinforce each other.

2.1.1 Prima Fascia Evidence for Status Quo Bias

The term 'status quo bias' was first coined in Samuelson and Zeckhauser [1988] in a paper of the same name. Their demonstrated the presence of such an effect in two different ways. First, they presented experimental subjects with choices that were framed in such a way as to make one alternative the status quo. These experiments showed strong and robust evidence of status quo bias – people were more likely to select a particular option when it was framed as the status quo than when it was not. Moreover, these effects occurred in choices between disparate types of objects – investment portfolios, paint color, cars and so on.

A potential criticism of any experimental study is that the results do not generalize to important 'real world' decisions. Samuelson and Zeckhauser therefore looked for status quo in another way, by examining the choices of health insurance plan by faculty members at their university. In particular, they were interested in a change implemented by the university that increased the number or health insurance plans that faculty members could choose between. Samuelson and Zeckhauser compared the choices of faulty members that joined the university before and after the increase in plans offered. What they found was that faculty members who joined before the increase were less likely to be enrolled in one of the new plans than those that joined afterwards. Assuming that these two groups were otherwise identical, this difference can only be explained by status quo bias on the part of those who joined before the change.

Following the work of Samuelson and Zeckhauser, there have been numerous other studies that have shown status quo bias in important economic decisions. Perhaps the most well known of these is a sequence of papers that look at the effect of automatic enrollment in 401(k) retirement plans (see, for example, Madrian and Shea [2000, 2001], Choi et al [2004], Beshears et al. [2009]).

These studies compare retirement savings behavior of people covered by 'opt-in' and 'opt-out' 401k retirement savings plans. The former refers to plans in which, if an employee does nothing they will not participate the 401k plan, while the latter refers to those in which employees are automatically enrolled in the plan unless they actively opt-out. Automatic enrollment means that employees who do nothing will have a default amount of their salary invested in a default investment fund. These studies find that opt-in plans not only have a huge effect on initial enrollment, they also affect the type of assets that people invest in and their savings rate. People covered by opt-in plans are much more likely to initially participate in the 401k plan, and to invest in the default fund and at the default rate, than those in opt-out plans. These differences diminish over time, but remain significant to roughly a 2 year time horizon.

Other studies have found evidence of status quo bias in people's choice of investment portfolio [Ameriks and Zeldes, 2001; Agnew, Balduzzi, and Sunden , 2003], mutual fund investment [Patel, Zeckhauser, and Hendricks, 1991; 1994]; Kempf and Ruenzi 2006], choice of electrical service provider [Hartman et al., 1991] and car insurance [Johnson et al., 1993]. Moreover, the presence of status quo bias is now taken into account in policy decisions, both at the governmental and firm level (see for example the Pension Protection Act [2006] and Benartzi and Thaler [2004])

2.1.2 Causes of Status Quo Bias: Theory and Evidence

While the research above gives ample evidence of the existence of status quo bias, it gives little understanding of what causes such a bias. In this section, I discuss some of the more common explanations for status quo bias, and the theoretical and experimental evidence for each cause.

Transaction Costs The most obvious explanation for status quo bias is that it is not a bias at all. In the presence of transaction costs – or a cost from changing for one alternative to another, it may be perfectly rational to stick with the alternative you currently have, even if you are offered a better alternative. For example, if it costs my $100 to change my cell phone provider, I may stick with my 'status quo', even if I think another provider is actually better. It is both intuitively plausible and empirically verified that transaction costs do affect the choices that people make.

Perhaps the key thing about transaction costs is that they cannot explain all the status quo bias that we observe, if they are defined only as the monetary cost or effort needed to switch from

one alternative to another. In many of the cases that I report above, transaction costs are low or non-existent. In Madrian and Shea [2001] the cost of switching to an alternative 401k plan is very low – a simple telephone call. In the experiments of Samuelson and Zeckhauser [1988], transaction costs are effectively zero. Thus, transaction costs are sufficient but not necessary for status quo bias.

The Endowment Effect One of the earliest explanations for status quo bias is a psychological phenomenon known as the endowment effect. This refers to the fact that people simply seem to like an object more because it is something they own. This effect was first demonstrated in a sequence of experiments by Knetsch and Sinden [1984], Coursey et al. [1987], Knetsch [1989] and Kahneman, Knetsch and Thaler [1990]. The classic experiment compares the behavior of two groups of subjects. The first group was given a bar of chocolate at the start of the experiment, and then later asked if they would like to trade it for a coffee mug. The second group were initially given the coffee mug, and asked if they would like to trade it for the chocolate bar. While about 80% of the first group chose the chocolate bar over the coffee mug, about 80% of the second chose the mug over the chocolate bar. Again, this is an environment in which the transaction costs were very low, suggesting that cannot be the explanation for what is going on.

The endowment effect has also been measured in another way – the gap between how much people would pay for an object and how much they would accept to sell the same object, were they to own it. This gap between willingness-to-accept and willingness-to-pay has been estimated experimentally, and can be very large. In an experiment performed using basketball tickets for Duke University games, Cameron and Ariely [2000] finds a willingness to pay of $170, and a willingness to accept of $2400, suggesting that those who owned the tickets valued them 14 times more than those that didn't. While these numbers should be treated with caution, they do suggest that the endowment effect can be large.

Some researchers see the endowment effect as a manifestation of loss aversion (see, for example, Tversky and Kahneman [1991]). This refers to the fact that people seem to weigh losses more heavily than gains. Thus, the loss of a mug due to selling it seems more important than the gain one gets from buying it. A person who suffered from loss aversion would exhibit the endowment effect.

It should be noted that more recent literature has highlighted circumstances under which peo-

ple do not exhibit the endowment effect [List, 2003; Plott and Zelner, 2005]. Perhaps the most important finding is that market experience seems to eradicate the endowment effect – people who are experienced in trading a particular item exhibit much reduced endowment effect.

Procrastination and Decision Avoidance A sizeable body of literature within the psychology profession documents the desire of people to avoid difficult or complicated decisions (see Anderson [2003] for a review). This decision avoidance can take many forms, but includes both the selection of the status quo (in order to avoid having to make an active choice) and procrastination (and therefore sticking with the status quo by default).

There is evidence that the difficulty of a choice task increases status quo bias. Tversky and Shafir [1992] and Shafir, Simonsen and Tversky [1993] demonstrate experimentally that people are more likely to stick with a status quo in decisions where the alternatives are difficult to compare. Iyengar and Lepper show that the number of alternatives between which a decision maker has to choose can effect status quo bias. The show that people are less likely to buy a pot of jam when asked to choose from a large choice set than when asked to choose from a small choice set – in other words they are more likely to stick with the status quo of 'no jam' in large choice sets than in smaller ones. Kempf and Ruenzi {2006] and Iyengar and Kamenica [2008] show similar choice set size effects in portfolio choice and 401k investments.

Perceived Recommendation A further mechanism by which a status quo, or default, option may affect choice is if it is perceived as a recommendation on the basis of whomever it is that decides on the status quo. For example, in opt-out 401k investment plans, the default investment fund may be interpreted as a recommendation on behalf of the employers. Indeed, it is usually the case that these funds are not chosen randomly, so this may be a sensible thing for a decision maker to do.

Evidence for the perceived recommendation effect comers from Madrian and Shea's [2001]analysis of 401k investment funds. They found that the default investment fund became significantly more popular amongst employees who were not affected by the auto enrollment. In other words, even people for whom the default fund was not the status quo became more likely to choose that fund. One explanation for this would be if these employees were interpreting the default as a recommendation.

Ambiguity Aversion and Learning One thing that separates the status quo option from other alternatives is that the decision maker may already own it. If this is the case, they will have had experience of that alternative, and understand its properties better than other alternatives. For example, if a decision maker's status quo browser is IE, then they may know its properties better than they do, for example, Chrome or Firefox. This in turn may lead the consumer to prefer their status quo. One reason is that doing so means that they do not have to make the effort of learning about other alternatives. Another is that it may be impossible to learn everything about other alternatives they have not owned. Choosing such alternatives would lead to ambiguous outcomes, which people tend not to like.

There is little direct evidence on learning and ambiguity aversion as a source of status quo bias, but there is ample evidence that people are, in general , ambiguity averse. This goes back to Ellsberg [1961]. It therefore seems likely that ambiguity aversion is an important contributor to status quo bias.

Inertia A final way of thinking about status quo bias is provided by Mandler [2004]. His proposal rests on two ideas – first that people need a good reason to move away from the current status quo, and second that people may not be able to easily compare all the alternatives that they are asked to choose between. These two hypothesis can generate status quo bias, because people will stick with the status quo until an alternative that is 'obviously better' comes along, which may happen only rarely when people find it difficult to compare evidence

Again, there is little direct evidence that tries to tease apart the effect of inertia from other explanations for status quo bias. However, inertia would imply the status quo would be more prevalent in more complicated choice settings (where people find it more difficult to compare alternatives). As I have previously discussed, there is evidence for such an effect.

3 Violations of 'Standard' Utility Functions

We now move onto a set of observations that, while do not violate the idea that people have a stable, well behaved utility function, do imply that this function may not be in line with the standard assumptions that economists usually make.

3.1 Social Preferences

Consider the following game:

Player A has $10 to divide between themselves player B. They can make a single 'take it or leave it offer' (for example they keep $7 and player B gets $3). Player B can then either accept the offer, and reject it, in which case both players get 0.

First of all, what is the Nash Equilibrium of this game? As this is a sequential game, we can solve this using backwards induction. In the last stage, player B has a choice between $0, and whatever player A offers them. They should therefore accept ANY offer, and player A should offer them 1 cent.

Is this what happens when this game (called the ultimaitum game) is played in the lab? Of course not. The fplayer A ususally offers player B between 50% and 20% of the pie, and player B usually rejects offers than less than 20%. This result comes as no surprise to anyone outside the field of economics, but it is a result that is NOT well captured by the standard models within economics. Though note that this is not necessarily a violation of the predictions of game theory, more a violation of the way in which we write down people's payoffs.

So what is going on here? There are three basic classes of explanation

1. Fairness: People have a preference for 'fair' outcomes over unfair ones
2. Altruism: People get utility from other people's happiness
3. Reciprocity: People like to punish people who have not treated them well.

Note that all these explanations have subtly different predictions. For example, reciprocity suggests that the reason player A gives player B more than 1 cent is that they beleive that player B will punish them by rejecting the offer if this is what they do. In order to test this, experimenters have worked with a game called the dictaitor game in which player B does not get to make a choice: they have to simply take whatever it is that A offers them. In such games it STILL seems to be the case that player A offers more that 0, apparently ruling out reciprocity as an explanation.

In my opinion, social preferences are a good example of a phenomena that is obviously important (and missing from standard models), but as yet we have not fully nailed down the best way to treat

it. In fact, this is currently an extremely active area of research in behavioral economics

3.2 Time Inconsistency

One topic that we have not covered in this course is how people make choices over time: For example: how do we choose how much to consume in any given period, or how much to save and so on and so forth. The standard way to do this in economics is to assume that people work out how much utility they would get from consumption in each period, then discount future utility. So, for example, the utility of consuming c_1 this period c_2 next period, c_3 the period after that and c_4 in the last period is

$$U(c_1, c_2, c_3, c_4) = u(c_1) + \beta u(c_2) + \beta^2 u(c_3) + \beta^3 u(c_4)$$

The key thing about this way of describing choice over time is that is is *time consistent.* This concept is actually quite subtle, but we can illustrate it in the following way: Imagine you were offered \$100 today or \$105 dollars next period, and you preferred \$100 today. Then it must also be the case that if I offered you \$100 in period 3 or \$105 in period 4, then you would prefer \$100 in period three. To see this, note that the first choice is between

$$u(100) \text{ and } \beta u(105)$$

while the second choice is between

$$\beta^2 u(100) \text{ and } \beta^3 u(105)$$

So, if $u(100) > u(105)$, then it must be the case that $\beta^2 u(100) > \beta^3 u(105)$. The reason that this is important is that it mean's people will not want to change their plans over time: If I make a plan to not consume 100 in period 3 but instead consume 105 in period 4, by the time I actually GET to period 3 then I will still be happy to save.

Unfortunately, it seems that people do not exhibit this sort of time consistency. Instead they suffer from temptation and self control issues. Imagine that I offered you a chouce between a cheesburger and a salad today. What would you choose? How about if I offered you the choice today, but you arn't going to eat whatever you choose until next week. What would you choose? Many people, it turns out, choose the cheeseburger today, but pick the salad for themselves in the

future. This time inconistency is important because it opens up the possiblity of people suffering from temptation, meaning that people may sometimes prefer to restrict their choice set (for example going to restaurants that do not serve cheeseburgers

3.3 Ambiguity

The final topic that we will talk about is ambiguity aversion. Again, ambiguity is not something that we have not had time to cover in detail in the lectures, but the idea is simple enough. Imagine that you are deciding whether or not to bet on a horse in a race. Unlike the case we dealt with in class (sometimes called risk) there is not an explicit probability given the horse will win - that is something that you have to figure out for yourself (this is why we describe it as ambiguity, rather than risk). The standard model of decision making in this case is that the decision maker figures out the probability they assign to the horse winning the race, then maximize their expected utility given this probability.

In fact, it turns out that people seem not to behave in this way. Instead they seem to dislike ambiguity, and avoid choices that involve an ambiguous option. This was first demonstrated in the famous Ellsberg paradox. The paradox is based on the following situation: I have filled an urn with 90 balls. 30 of those balls are red, and 60 are either black or yellow. Which of the following bets would you prefer?

Consider the following pairs of choices

A You win $100 if you draw a red or a yellow ball

B You win $100 If you draw a black or a yellow ball

and

C You win $100 if you draw a red ball

D You win $100 if you draw a black ball

Most people faced with these choices choose C over D and B over A. However, this is not consistent with the standard model, which (basically) assumes that you should guess a number of

black balls and act accordingly. However, if this were the case, the fact that you chose C over D indicates that you assume that there are more red balls than black balls (i.e. more than 30). This should imply that you think that there are more than 30 yellow balls, and so you should prefer B over A. In fact, people always choose the option with no ambiguity - acting as if there are less than 60 yellow and black balls in total/.