Behavioral Economics: An overview

By Lucas Engelhardt

1 Introduction

The standard neoclassical model of behavior makes a few simple assumptions about human behavior and preferences:

1. Preferences are stable.
2. Preferences are transitive.
3. Preferences are individualistic.
4. More is preferred to less (for a good).
5. Preferences are over outcomes.
6. Goods display diminishing marginal rates of substitution.
7. People are reasonably good at statistical reasoning.

Behavioral economics sets out to study these assumptions somewhat directly, and find that many of these assumptions are often not true.

2 Endowment Effects and Loss Aversion

Suppose that mugs cost $2 a piece, and you have an income of $20. In that case, you have a standard budget line like in this diagram: One important result from the standard theory is that it doesn't matter where on the budget line you start, you should end up at the same place. So, if I give you $20, you should end up in the same state as if I give you 10 mugs. The budget line is the same, the indifference curves are the same, and therefore the choice should be the same.

However, behavioral economists have found some evidence against this claim. To test it, they performed the following experiment:

First, they divided a group of people into two groups: one received a mug was asked how much they'd be willing to accept to give the mug up, the other was not given a mug and was asked how much they'd be willing to pay for a mug. According to the standard theory, on average the two groups should give you the same number. Why so?

Autumn 2009/Mugs.jpg

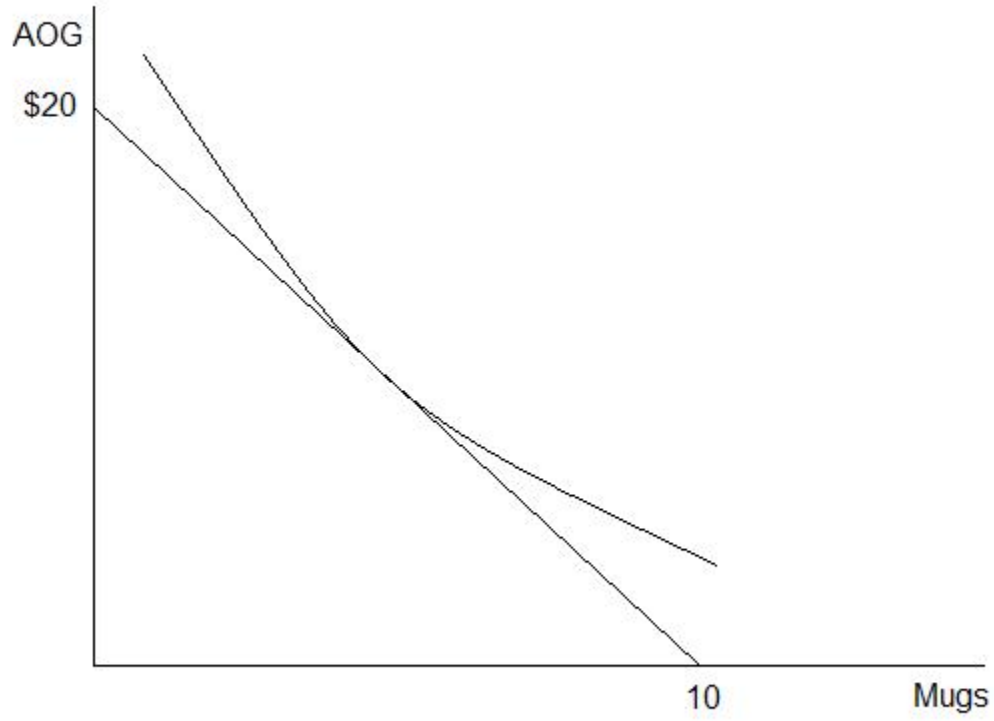

We can think of this problem as trying to narrow down the shape of indifference curves. If a person is just willing to accept a certain price for the mug, it must be that they are nearly indifferent between having either that much money or the mug. If someone is just willing to pay a certain price for the mug, it must be that they are nearly indifferent between having either that much money or the mug. Since the people were divided randomly, there's little reason to believe that the two groups will have systematically different preferences unless preferences depend on whether or not you start out with a mug or without one. Under standard assumptions, preferences don't depend on where you start on a budget line - so it doesn't matter whether you started with a mug or not. That is, an indifference curve should look something like this on average: So, it doesn't matter whether you start at 0 mug and are asked about your willingness to pay, or whether you start at 1 mug and are asked about your willingness to accept - the answer should be the same.

Kahneman, Jack L. Knetsch, and Thaler (1990) tried this experiment, and found that the median willingness to accept was $5.75. The median willingness to pay was $2.25. Not only is there a difference between

Autumn 2009/WTPWTA.jpg

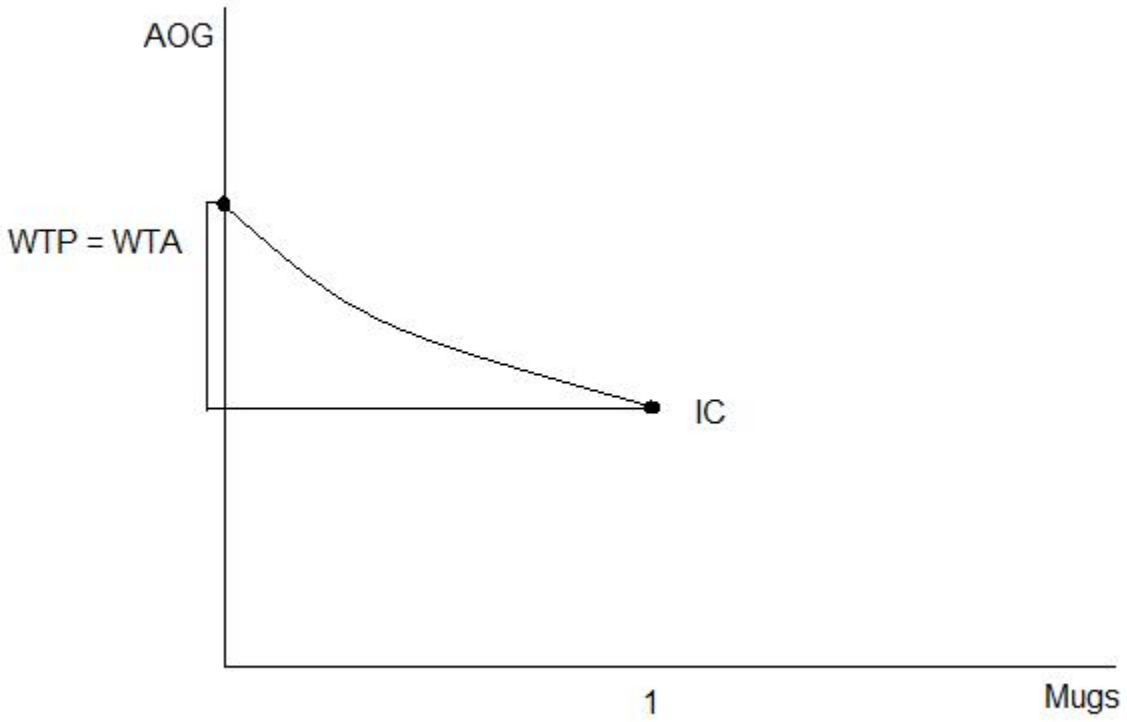

the two, but there is a very large difference! Once given the mug, people seem to be very hesitant to part with it, but if they aren't given the mug in the first place, they don't seem particularly inclined to get it. This is called the "endowment effect". People tend to value something more after they have it than before. (So, it ends up, the grass *isn't* greener on the other side of the fence!) This implies "loss aversion" - that is, people tend to try to avoid losing what they already have, and are more willing to put forth effort to avoid loss than to gain the thing in the first place.

A careful student will object: "But the choice is not the same! The budget line for those that received the mug is higher!" This is absolutely true. At the beginning of the experiment, those who are given the mugs are, in fact, given the mugs in addition to any other wealth that they already have. So, it may be that people who are about "one mug poorer" are less willing to pay for the mugs.

To fix that problem, Kahneman, Knetsch, and Thaler ran another experiment. One group was given the mug (as before), the other was given a choice between either a mug, or a certain sum of money. That sum was varied so that they could approximate indifference. The results of this experiment? The median person who was given the mug were willing to accept $7.12. The median person who was given a choice between a mug and money valued the mug at about $3.12. So, even when we control for wealth effects, it seems that people still show an "endowment effect".

However, that's not the end of the story. Some have suggested that this experiment relies heavily on two facts: (1) lack of anonymity. Perhaps people want to have a reputation of "driving a hard bargain". So, sellers of a good should report that they are only willing to accept a very high price, and buyers should report that they are only willing to pay a very low price. (2) lack of experience. When people aren't comfortable, they may display an endowment effect only because they are avoiding thinking about the alternative (since thinking is hard, after all!) - and are sticking to the default more than they actually might if they didn't have to think so hard.

To control for these effects List (2003) ran an experiment at a sports card fair. He assigned sports memorabilia A or B as compensation for filling out a questionnaire. Then, the recipient was allowed to switch to the other choice, if they so desired. The two items were of equal value, so, if there was no endowment effect, about 50 percent of recipients should have switched. List found that, for relatively inexperienced traders, only 6.8% switched. So, there is significant evidence of an endowment effect for this group. However, for very experienced traders (who trade at least six times a month), 46.7% traded - so there is virtually no endowment effect for this group. List also performed experiments eliciting willingness to pay and willingness to accept for different levels of experience. His findings were consistent with what he found before: inexperienced traders have a very large endowment effect; experienced traders do not.

I'm sure you've thought of the objection. Isn't it likely that people who naturally have small endowment

effects will trade more often? If you're someone who has a very small level of loss aversion, your willingness to pay and willingness to accept will be very close together, so lots of trades will look good. So, this may be a case where there's an "omitted variable". Some people are just naturally more loss averse, and will therefore be unlikely to trade much, and unlikely to trade the prize here. Some people are just naturally less loss averse, and will therefore be likley to trade a lot, and likley to trade the prize here. So, while this may establish that at least some people are loss averse, it doesn't establish that inexperience in trading causes *causes* loss aversion.

To answer that objection, List ran another experiment where he kept track of the same traders over several months. What he found: those traders that increased their trading frequency also decreased their loss aversion, and those that increased trading frequency the most displayed the biggest decreases in endowment effects. This result has a couple interpretations. One possibility is that frequent trading makes it more obvious to the trader that he is acting in a loss-averse manner - and therefore the frequent trader tries to act differently. Another possibility is that the "endowment" changes. A frequent trader expects that they will trade the good, and so, even when it is assigned, it is not considered part of the trader's "endowment" in the same sense as before.

3 Reference Points, Backward Sloping Supply of Labor, Narrow Framing, and Anchoring

A related topic is about "reference points". The suggestion in the "reference point" hypothesis is that people compare a state of affairs to some "reference point", and, if things are "worse", then that is considered to be "very, very bad". If things are "better", then things are considered to be "good" (but not "very, very good"). The result, then, is that people will work very hard to attain the reference point, but not very hard to attain anything better. (This is similar to the endowment effect and loss aversion - people who are given a mug have "having a mug" as their reference point, and require large compensation to give it up.)

This type of behavior should be familiar to students. While there are some students who want to learn as much as possible in every class, these students are generally "few and far between". Most students have specialized interests where they really do want to learn as much as possible, but are forced to take some classes outside their interests as well. It's not unusual for a student taking a class outside of their interests to have some minimum level they are trying to attain, to put forth a lot of effort to get to that level, but to put forth almost no effort after they attain it. ("I just need a C in this class!")

Similar behavior can be found in labor markets.

The standard story for labor supply is this: As the wage increases, the "price of income" (in terms of leisure) decreases. This creates a substitution effect where people consume more income and less leisure. It also creates an income effect where people consume more leisure (since leisure is a normal good). If the substitution effect is larger than the income effect, then labor supply slopes up - as the wage increases, people work more. If the income effect is larger than the substitution effect, then labor supply slopes down - as the wage increases, people work less. Typically, economists believe that, for low wages, the substitution effect dominates, but for high wages the income effect dominates. So, you get a "sideways U" shape like this: The

Autumn 2009/LaborSupply.jpg

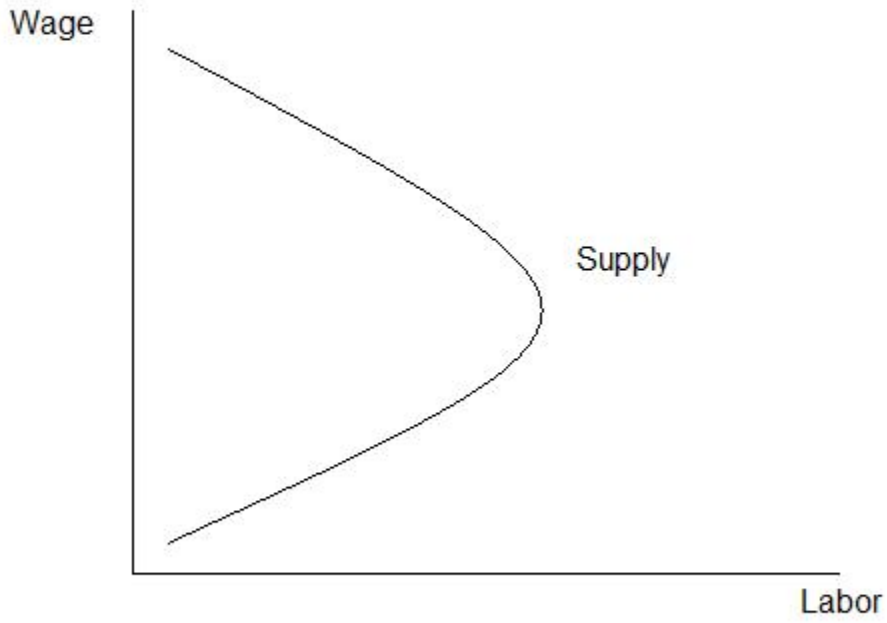

reference point hypothesis makes a different claim.

According to the reference point hypothesis, a worker is trying to earn a particular total income, so that "wages" times "hours worked" should be some constant. If this were strictly true, then labor supply would have this shape: An odd prediction: as wages drop, people work more and more and more... if wages are sufficiently low, people will work 24 hours a day.

Autumn 2009/LaborSupply2.jpg

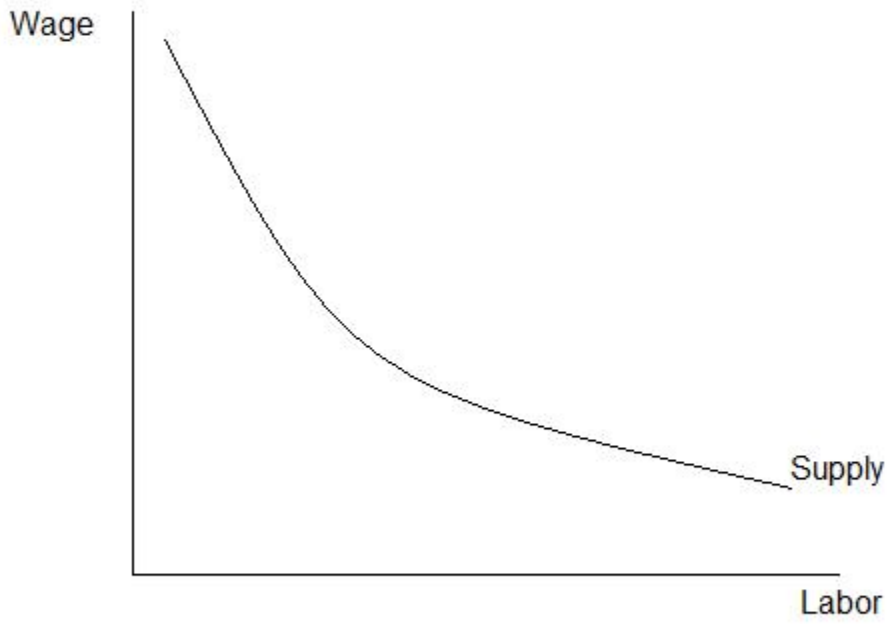

Reference point theorists suggest that this is not strictly true. People do have a target income in mind, but are willing to deviate from it if the benefit is great enough. For example, working 20 hours a day to attain the target income is so hard that workers will not do that. So, below a certain wage, labor supply is upward sloping. This doesn't look too different from the standard theory. However, reference point theory doesn't stop here! It also claims that, if the wage is high enough, people are actually willing to work more to earn more than the target wage! So, above a certain wage level, labor supply is upward sloping AGAIN. The end result? Of course, an attentive student should be immediately suspicious of this model. After all, people generally don't get to observe their wage rate and then decide how many hours to work. Most people are offered job choices like "10 hours a week for $10 an hour" as opposed to "30 hours a week for $20 an hour", where the hours/wage pair is a pair.

However, there are exceptions to this rule - and those exceptions allow for testing of this hypothesis. Camerer et al. (1997) use data from New York cab drivers - who do get to decide how long they work.

Autumn 2009/LaborSupply3.jpg

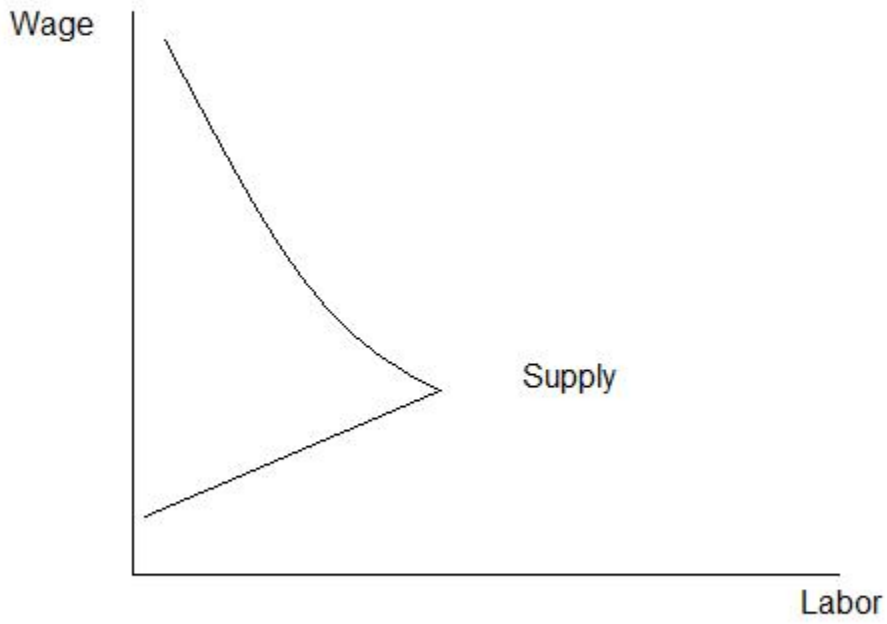

While they can't test the "S" shape, they can test whether supply is upward or downward sloping. Since cab drivers aren't "high wage" by most standards, the standard model would put them on the upward sloping portion of the supply curve. However, it is possible that cab drivers have a target income such that the wage they do earn is on the downward sloping part of the "S" shaped reference-point based curve. Here's what they find: cab drivers tend to work longer when wages are low. At first glance, this seems to support a downward sloping supply curve. But, this is not necessarily so. It may be that the supply curve is upward sloping, but that the supply curve shifted between observation points. (For example, the fact that I work 6 hours and earn \$15 an hour one day and work 7 hours and earn \$12.50 might just reflect the fact that on the first day it was raining, so there were fewer cabbies on the road. So, the data points are actually points on a *demand* curve, not on a supply curve.) To solve this problem, Camerer et al. used various techniques to hold the supply curve "constant", and find that the results still indicate a downward sloping demand curve. Some economists suggest that their techniques are still not good enough.

Autumn 2009/LaborSupply4.jpg

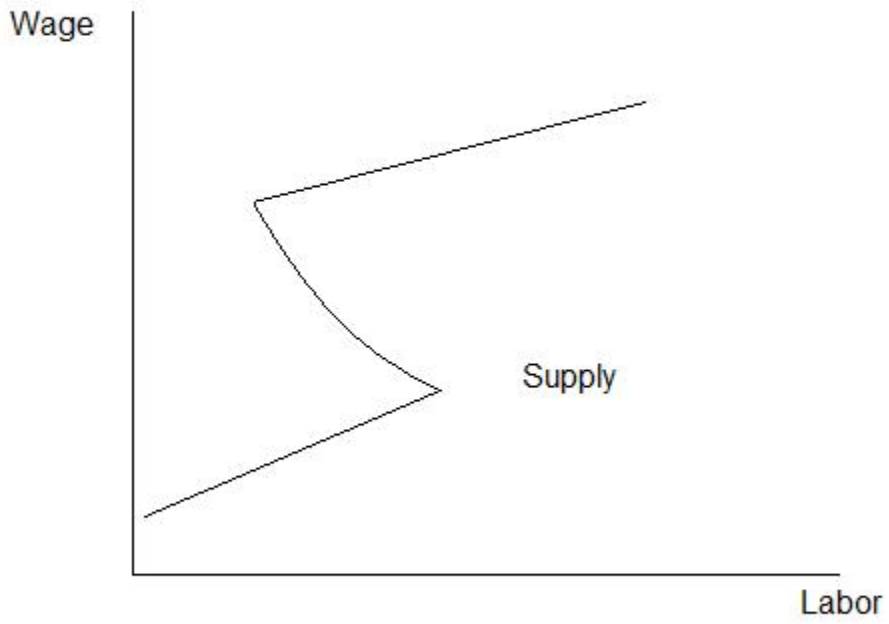

To avoid these problems Fehr and Goette (2007) did an experiment to control for "supply curve" effects. They divided a group of bicycle messengers into two groups. Each group was given a 25% higher commission on their deliveries in one of two months (they were told which month they would get a higher commission in). What they found: those in the "high wage month" signed up for 30% more shifts, and did 6% fewer deliveries within a shift. The first fact is consistent with both the reference point model and the standard model. In the standard model, the higher wage induces working more. In the reference point model, messengers will choose to work during those days when it is easiest to meet the target income. The second fact is not consistent with the standard model, but is consistent with the reference point model. This suggests that reference points may be important in labor supply decisions, at least for some workers.

A related concept is "narrow framing". People aren't very good at thinking across categories of goods. As an example: Suppose that you're planning on going to see a show. The tickets are $10. Consider two cases: In the first scenario, you're planning to buy the ticket at the door, but when you get to the door,

you find that $10 is missing from your wallet. You have plenty of money to buy another ticket. Do you buy the ticket? In the second scenario, you've already purchased the ticket, but when you get to the door, you find that the ticket you had purchased is missing. There are plenty of tickets for sale, and you have enough money in your wallet to buy a ticket. Do you buy it? For most people, in the first case, the answer is "yes", but in the second case, the answer is "no". However, from an economist's perspective, the two situation are identical. In both cases, you have exactly the same amount of money - therefore the same budget. In both cases, you expected to see the show - so your expectations were identical. It's reasonable that the resulting choice should be identical. Yet, for most people it isn't. Why not? Because of "narrow framing". People don't think of the ticket as being "equivalent" to $10. Similarly, people are willing to drive across town to save $0.10 a gallon on 10 gallons of gas (total: $1), but aren't willing to drive across town to save $1 on a television set. This is another example of narrow framing. Since $0.10 is a not insignificant part of the price of gasoline, people perceive it as a "big" savings - while the $1 is a small part of the price of a television set. Rather than thinking that "a dollar is a dollar", people seem to weight dollars based on their purpose, seemingly forgetting that a dollar saved in one way is exactly as "spendable" as a dollar saved in any other way!

One last application of the "reference point" idea is "anchoring". People seem influenced by information that is totally unrelated to the problem at hand. For example, a set of subjects was asked what the last two digits of their social security number were. Then, they were asked if they were willing to purchase some item for that price (the item was worth about $70). Finally, they were asked the maximum price they were willing to pay for the item. Subjects with higher social security numbers systematically reported higher prices. Is it true that people with high social security numbers are just "big spenders"? Probably not. Social Security numbers are assigned nearly randomly. So, a more likely explanation is that calling the numbers to mind gave the mind a "reference point" - for an entirely unrelated problem. (Note: notice how many commercials say "Would you be willing to pay X dollars? Y dollars? Well, we're only charging Z dollars!" There is often no evidence that the "X" and "Y" prices have any meaning whatsoever - but they may help establish a reference point in the mind of the buyer.)

4 Gift Exchange

The standard model claims that people are self-interested. However, there is some data to suggest that people often display "social preferences", especially in cases where they believe that other poeple also have "social preferences". (A case of "Since you scratched my back, I'll scratch yours - even though no strings were attached.")

One classic case of this is "Gift Exchange". If one person gives another a gift, the other person often reciprocates by giving a gift back. One application of this is in labor markets. Gift exchange models suggest that a worker will respond to a higher wage by working harder. One paper that suggests that this is true is Kube, Marechal and Puppe (2008).

They performed an experiment in which people were hired to work for three hours shelving books in a library. They were each hired for a total of 36 euros for their time. The hirees were then divided into three groups. One group received the agreed wage. The second group received the agreed wage, plus a 20% (7 euro) bonus (which they were told about before they started working). A third group was given a thermos that was worth about 7 euro.

The results are somewhat surprising. The group that received a cash bonus was about 6% more productive than the group that received no bonus. So, the cash gift did increase productivity - but it did not pay for itself (a 20% increase in costs for a 6% increase in productivity). The group that received the thermos however, was 30% more productive than the group that received no bonus. So, the thermos gift did pay for itself.

This is quite suggestive of "Gift Exchange" motives. If people perceive a wage increase as simply a wage increase, then it is unlikely to create "Gift Exchange" reactions. Therefore, the effect of cash gifts is relatively small, as they are perceived as a wage increase. On the other hand, gifts "in kind" are perceived as gifts, and, as a result, induce large "exchange" effects.

To be safe, though, we have to rule out the possibility that this thermos might just be worth more than the cash. No worries! In a separate experiment, people were given the choice between 7 euro and the thermos - and they overwhelmingly chose the money. So, the effect from the experiment isn't just that the thermos was seen as a "high value" gift - the effect was from the fact that it was definitely a gift, rather than "payment for services rendered".

5 Violations of Transitivity

One common tendency that behavioral economists find: Preferences are often not even transitive!

How this was tested: People were asked about how much they'd be willing to pay to be party to a couple gambles. One gamble was a "high stakes, low probability gamble". For example, they had a 1% chance of winning $1,000. The other gamble was a "low stakes, high probability gamble". For example, they had a 90% chance of winning $11. Most people will assign a relatively high value to the "high stakes" gamble - this seems to indicate that they prefer the high stakes gamble - yet many people, when given the choice between the two gambles, will take the low stakes gamble! This is a violation of transitivity. Based on valuation, we know that, for example: $6.01 ¿ "high stakes gamble" ¿ $6. That is, the gambler is willing to pay $6, but

not \$6.01 for the high stakes gamble. Also, \$5.01 ¿ "low stakes gamble" ¿ \$5. That is, the gambler is willing to pay \$5, but not \$5.01 for the high stakes gamble. If we make the reasonable assumption that six dollars is prefered to five dollars and one cent, then it follows that the high stakes gamble should be preferred to the low stakes gamble - but that's not what a sizeable minority of people do! In one experiment, about 1/3 of the subjects assigned a higher value to the high stakes gamble, but then chose the low stakes gamble. This suggests that, perhaps, preferences aren't even transitive!

6 Failures in Statistical Reasoning

The fact is that most people are very bad at statistics. There are a few trends in the "badness".

First, overconfidence. This tends to come in three types. The first is overconfidence in ability. In a class with a curve so that 30% of students get A's, the grand majority will expect to get A's - or will at least overestimate the odds of getting an A. Most people overestimate their own abilities. In a survey, Svenson (1980) found that 93% of drivers rate their driving ability as "above the median". People consistently underestimate the odds of certain negative events like hospitalization and the time needed to finish a project. Camerer and Lovallo (1999) ran a multiple round experiment in which the top "c" out of "n" entrants in each round received positive profits (players are allowed to not enter - guaranteeing a payoff of zero). In one set of rounds, the "top" players were determined by luck. In another set of rounds, the "top" players were determined by their ability to solve a particular puzzle. Far more players entered in the second set, indicating that the average player overestimated their own ability, relative to the abilities of others. Similar behavior is seen in the field when people seem over confident in their ability to exercise self-control. So, people will join health clubs and never go (for example). If people knew that they would never go, it makes little sense for them to join. However if people overestimate their likelihood of going to the gym, then they may pay for a membership.

Second, the "law of small numbers". People often take small, but well known, events and overestimate their probability. For example, after 9/11, people consistently overestimated the odds of planes being hijacked - despite the fact that new security procedures decreased that probability. It is quite common that people make this type of mistake. A single story (success or horror!) can have a huge impact on people's expectations - even relative to good statistical work!

7 Conclusions

Naturally, there's a lot more out there about Behavioral Economics that I can't cover here. I'd highly recommend DellaVigna's recent article in the Journal of Economic Literature - "Psychology and Economics: Evidence from the Field". It's a bit technical on some point, but the basic ideas definitely come through. Much of this piece is from the evidence catalogued by DellaVigna. There's also a classic piece by Kahneman and Tversky that started much of the thinking about this field. It's titled "Judgment Under Uncertainty: Heuristics and Biases" - it appeared in Science in 1974.

So what is the main "gist" of all this?

(1) Standard assumptions are likely not to be true. (2) We should be somewhat cautious in utilizing models that require too much information processing of too many agents - odds are that something is likely to go wrong. (3) A firm understanding of some of the behavioral characteristics may help you make better decisions - either because you counteract their effects, or because you take advantage of them!