Behavioural Economics 101:

Judgment, Choice and Time

Behavioural Economics increases the explanatory power of Economics by providing it with realistic psychological foundations. A pioneer of the field explains.

by George Loewenstein

WHEN ECONOMICS WAS FIRST IDENTIFIED as a distinct field of study in the late 19th century, Psychology did not yet officially exist as a discipline. Nevertheless, many economists moonlighted as the psychologists of their times.

Adam Smith, best known for *The Wealth of Nations*, wrote a less well-known book, *The Theory of Moral Sentiments*, which laid out psychological principles of individual behaviour that are arguably as profound as his economic observations. The book is bursting with insights about human psychology — many of which presage current developments in Behavioural Economics. For example, Smith commented that "We suffer more... when we fall from a better to a worse situation, than we ever enjoy when we rise from a worse to a better." Loss aversion!

Research in the field of behavioural decision research — on which Behavioural Economics draws more than any other subfield of Psychology — typically falls into two categories: *judgment* and *choice*. Judgment research deals with the processes people use to estimate probabilities, while choice research deals with the processes people use to select among actions, taking account of any relevant judgments they may have made.

In this article, I will focus on the key findings from these two foundational aspects of Behavioural Economics and discuss some of the field's most promising developments.

Judgment Research: Key Findings

Will the Fed raise interest rates? Will you lose your job in a downturn? Will you be able to find another house you like as much as the one you must bid for right away? Will it rain during your vacation to London? Judging the likelihood of events is central to economic life.

Cognitive psychologists have proposed several 'heuristic' mechanisms that lead to judgments that sometimes violate pure rationality. For example, people often judge the probabilities of future events based on how easy those events are to imagine or to retrieve from memory. This *availability heuristic* contributes to many further biases. One is *hindsight bias*: Because events which

A good heuristic provides fast, close-to-optimal answers when time or cognitive capabilities are limited; but it can also lead to errors.

actually occurred are easier to imagine than counterfactual events that did not, people often over-estimate the probability they previously attached to events that later happened. This bias leads to 'second-guessing' or 'Monday-morning quarterbacking', and may be partly responsible for lawsuits against stockbrokers who have lost money for their clients (i.e. the clients think the brokers 'should have known').

A more general bias is the *curse of knowledge*: People who know a lot find it hard to imagine how little others know. The development psychologist **Jean Piaget** suggested that the difficulty of teaching is caused by this curse. Anybody who has tried to learn from a computer manual has seen the curse of knowledge in action. Another heuristic for making probability judgments is called *representativeness*: People judge conditional probabilities by how well the data represents the hypothesis or the example represents the class. Like most heuristics, *representativeness* is an economical shortcut that delivers reasonable judgments with minimal cognitive effort in many cases — but sometimes goofs badly and is undisciplined by normative principles. Prototypical exemplars of a class may be judged to be more likely than they truly are (unless the prototype's extremity is part of the prototype).

For example, in judging whether a certain student described in a profile is, say, a Psychology major or a Computer Science major, people instinctively dwell on how well the profile matches the Psychology or Computer Science major stereotype. Many studies show how this sort of 'feature-matching' can lead people to underweigh the base rate — in this example, the overall frequency of the two majors.

Another by-product of *representativeness* is the *law of small numbers*, whereby small samples are thought to represent the properties of the statistical process that generated them. If a baseball player gets hits 30 per cent of his times at bat, but is zero-for-four so far in a particular game, then he is 'due' for a hit in his next at bat in this game, so that this game's hitting profile will more closely represent his overall ability.

The so-called *gambler's fallacy* — whereby people expect a tail after a coin lands heads three times in a row, is one manifestation of the law of small numbers. The flip side of the same misjudgment (so to speak) is surprise at the long streaks which result if the time series is random, which can lead people to conclude that the coin must be unfair, when it isn't. Field and experimental studies with basketball shooting and betting on games show that people, including bettors, believe that there is positive autocorrelation — that players experience the 'hot hand' — when there is no empirical evidence that such an effect exists.

It is important to note that heuristics can be good or bad: A good heuristic provides fast, close-to-optimal answers when time or cognitive capabilities are limited; but in some situations, it also violates logical principles and leads to errors.

Choice Research: Key Findings

Standard preference theory incorporates a number of assumptions. For example, it assumes that preferences are 'reference independent' — i.e., they are not affected by an individual's current state or context. It also assumes that preferences are not influenced by variations in the way that options are presented.

Both of these assumptions have been disproven by behavioural researchers. For example, numerous *framing effects* show that the way choices are presented to an individual often determine the preferences that are then 'revealed'. The classic example is **Daniel Kahneman** and **Amos Tversky**'s Asian disease problem, in which people are informed about a disease that threatens 600 citizens and asked to choose between two undesirable options. In the 'positive frame', they are given a choice between (a) saving 200 lives for sure, or (b) a 1/3 chance of saving all 600 with a 2/3 chance of saving no one. In the 'negative frame', people are offered a choice between (c) 400 people dying for sure, or (d) a 2/3 chance of 600 dying and a 1/3 chance of no one dying. Despite the fact that A and C, and B and D are equivalent in terms of lives lost or at risk, most people choose A over B, and D over C.

Another phenomenon that violates standard theory is called an *anchoring effect*. The classic demonstration of this effect was identified in the context of judgment rather than choice. Subjects were shown the spin of a wheel of fortune that could range between 0 and 100 and were asked to guess whether the number

of African nations in the **United Nations** was greater than or less than this number. They were then asked to guess the true value. Although the wheel of fortune was obviously random, subjects' guesses were strongly influenced by the spin of the wheel. As Kahneman and Tversky interpreted it, subjects seemed to 'anchor' on the number spun on the wheel and then adjusted for whatever else they thought or knew, but adjusted insufficiently.

Anchoring effects have also been demonstrated for choices as opposed to judgments. In one study, subjects were asked whether their certainty equivalent for a gamble was greater than or less than a number chosen at random, and then were asked to specify their actual certainty equivalent for the gamble. Again, the stated values were correlated significantly with the random value.

Many studies have also shown that the method used to elicit preferences can have dramatic consequences, sometimes producing *preference reversals* — situations in which A is preferred to B under one method, but A is judged as inferior to B under a different elicitation method. The best known example contrasts how people choose between two bets versus what they separately state as their selling prices for the bets. If bet A offers a high probability of a small payoff and bet B offers a small probability of a high payoff, the standard finding is that people choose the more conservative bet A over bet B, but are willing to pay more for the riskier bet B when asked to price them separately.

Another form of *preference reversal* occurs between joint and separate evaluations of pairs of goods. People will often price or otherwise evaluate an item A higher than another item B when the two are evaluated independently, but evaluate B more highly than A when the two items are compared and priced at the same time. *Context effects* refer to ways in which preferences between options depend on what other options are in the set. For example, people are generally drawn to 'compromise' alternatives whose attribute values lie between those of other alternatives.

All of these findings suggest that preferences are not the pre-defined sets of indifference curves represented in Microeconomics textbooks: In reality they are often ill-defined, highly malleable and dependent upon the context in which they are elicited.

A theme emerging in the behavioural research is that, although people often reveal inconsistent or arbitrary preferences, they typically obey normative principles of economic theory when it is transparent how to do so. Researchers refer to this as *coherent arbitrariness* and illustrated the phenomenon with a series of studies in which the amount of money subjects demanded to listen to an annoying sound was sensitive to an arbitrary anchor (a random amount of money that was based on their social security number). Although the impact of this number revealed a degree of arbitrariness in subjects' valuations, subjects, sensibly, demanded much more to listen to the tone for a longer period of time. Thus, while expressed valuations for one unit of a good are sensitive to an anchor that is clearly arbitrary, people also obey the normative principle of adjusting those valuations to the quantity — in this case, the duration — of the annoying sound.

As indicated, most of the evidence that preferences are constructed comes from demonstrations that some contextual features that *should not matter* actually *do matter*. Whether it be the composition of a choice set, the way gambles are 'framed' as gains or losses from a reference outcome, or whether people choose among objects or value them separately, all have been shown to make a difference in expressed preference.

TIME DISCOUNTING. A subset of the choice research has become one of the key topics in Behavioural Economics: How individuals trade off costs and benefits that occur at different points in time.

The standard assumption is that people rationally weight future utilities by an exponentially-declining discount factor. **Richard Thaler** was the first to empirically test the constancy of discounting with human subjects. Prof. Thaler told subjects to imagine that they had won some money in a lottery held by their bank: They could take the money now or earn interest and wait until later. They were asked how much they would require to make waiting just as attractive as getting the money immediately.

Prof. Thaler then estimated implicit (per-period) discount rates for different money amounts and time delays under the assumption that subjects had linear utility functions. Discount rates declined linearly with the duration of the time delay.

Preferences are often ill-defined, highly malleable and dependent upon the context in which they are elicited.

Later studies replicated the basic finding that discount rates fall with duration.

The most striking effect was an *immediacy effect*: Discounting is most dramatic when one delays consumption that would otherwise be immediate. *Hyperbolic time discounting* implies that people will make relatively far-sighted decisions when planning in advance — when all costs and benefits will occur in the future — but will make relatively short-sighted decisions when some costs or benefits are immediate. The systematic changes in decisions produced by *hyperbolic time discounting* create a time-inconsistency in intertemporal choice not present in the exponential model: An agent who discounts utilities exponentially would, if faced with the same choice and the same information, make the same decision prospectively as he would when the time for a decision actually arrives. In contrast, somebody with *time-inconsistent hyperbolic discounting* will wish prospectively that in the future he will take far-sighted actions; but when the future arrives, he will behave against his earlier wishes, pursuing immediate gratification rather than long-run well-being.

Most big decisions in life — savings, educational investments, hiring, health and diet — have costs and benefits that occur at different points in time. Many authors have discussed the issues of self control and stressed their importance for economics. An important question in modelling self-control is whether agents are aware of their self-control problem ('sophisticated agents') or unaware ('naive agents'). There are certainly many times in which people are partially unaware of their own future misbehaviour, and hence overly optimistic that they will behave in the future the way that their 'current self' would like them to. Researchers have shown that awareness of self-control problems can powerfully moderate the behavioural consequences of quasi-hyperbolic discounting.

Naivete typically makes damage from poor self-control worse. For example, severe procrastination is a creation of over-optimism: One can put off doing a task repeatedly if the perceived costs of delay are small — 'I'll do it tomorrow; there is little loss from not doing it today' — and hence accumulate huge delay costs from postponing the task many times. A sophisticated agent aware of his procrastination will realize that if they put it off now they will put if off in the future, and hence will do the task immediately.

However, in some cases, being sophisticated about one's self-control problem can exacerbate yielding to temptation. If you are aware of your tendency to yield to a temptation in the future, you may conclude that you might as well yield now; by the same token, if you naively think you will resist temptation in the future, you might feel 'licensed' to indulge in the present.

An anomaly to this is *negative time discounting*. If people like savouring pleasant future activities, they may postpone them to prolong the pleasure (and they may also try to get painful activities over with quickly, to avoid dread).

The Rise of 'Behavioural Finance'

Until fairly recently, financial theory bet all of its chips on the belief that investors are so rational that any observed historical patterns that could be used to beat the market are detected — the 'Efficient Markets Hypothesis'. In 1978, **Michael Jensen** called this hypothesis "the most well-established regularity in social science." But, shortly after his grand pronouncement, the list of anomalies began to grow.

One important anomaly is the 'equity premium puzzle': Average returns to stocks are much higher than returns to bonds (presumably to compensate stockholders for higher perceived risks). To account for this pattern, **Shlomo Benartzi** and Richard Thaler assumed a combination of *decision isolation* — investors evaluate returns using a one-year horizon — and *aversion to losses*. These two ingredients create much more perceived risk to holding stocks than would be predicted by expected utility.

Another anomaly is the magnitude of volume in the market. The so-called '**Groucho Marx**' theorem states that people should not want to trade with people who would want to trade with them — but the volume of stock market transactions is staggering. For example, **Terrance Odean** has noted that the annual turnover rate of shares on the New York Stock Exchange is greater than 75 per cent, and the daily trading volume of foreign-exchange transactions in all currencies (including forwards, swaps and spot transactions) is equal to about one-quarter of the total annual world trade and investment flow. Odean then

presents data on individual trading behaviour which suggests that the extremely high volume may be driven, in part, by over-confidence on the part of investors.

Early heretics like Nobel Laureate **Robert Shiller**, who argued empirically that stock price swings are too volatile to reflect only news, and **Werner DeBondt** and Richard Thaler, who discovered an important over-reaction effect based on the psychology of representativeness, had their statistical work 'audited' with special scrutiny (or worse, were simply ignored). Today, a younger generation of academics and finance professionals is eagerly sponging up as much psychology as they can to help explain how efficient markets truly are.

In closing

Critics have pointed out that Behavioural Economics is not a unified theory, but instead, a collection of tools and ideas. This is true. It is also true of neoclassical Economics. A worker might rely on a 'single' tool — say, a power drill — but also use a wide range of drill bits to do various jobs. Is this one tool or many?

Likewise, economic models do not derive much predictive power from the single tool of utility-maximization. Precision comes from the 'drill bits' — such as time-additive separable utility in asset pricing including a child's utility into a parent's utility function to explain bequests, rationality of expectations for some applications and adaptive expectations for others, homothetic preferences for commodity bundles, price-taking in some markets and game-theoretic reasoning in others, and so forth.

Sometimes these specifications are even contradictory. For example, pure self-interest is abandoned in models of bequests, but restored in models of life-cycle savings; and risk-aversion is typically assumed in equity markets and risk-preference in betting markets. Such contradictions are like the 'contradiction' between a **Phillips**-head and a regular screwdriver: They are different tools for different jobs. The goal of Behavioural Economics is to develop better tools that, in some cases, can do both jobs at once.

In the end, all Economics rests on some sort of implicit psychology. The only question is whether the implicit psychology in Economics is good psychology or bad psychology. Given what we know to be true, it is simply unwise, and inefficient, to 'do' Economics without paying at least *some* attention to good psychology. ***RM***

George Loewenstein is the Herbert A. Simon University Professor of Economics and Psychology at Carnegie Mellon University, co-director of its Centre for Behavioural Decision Research, and Director of Behavioural Economics at the Centre for Health Incentives at the Leonard Davis Institute of the University of Pennsylvania. This article is an adaptation of his seminal chapter in *Advances in Behavioural Economics* (Princeton University Press, 2003) written with **Colin F. Camerer**, the Robert Kirby Professor of Behavioural Finance and Economics at the California Institute of Technology.