

Supposedly Irrelevant Factors

Early in my teaching career I managed to inadvertently get most of the students in my microeconomics class mad at me, and for once, it had nothing to do with anything I said in class. The problem was caused by a midterm exam.

I had composed an exam that was designed to distinguish among three broad groups of students: the stars who really mastered the material, the middle group who grasped the basic concepts, and the bottom group who just didn't get it. To successfully accomplish this task, the exam had to have some questions that only the top students would get right, which meant that the exam was hard. The exam succeeded in my goal—there was a wide dispersion of scores—but when the students got their results they were in an uproar. Their principal complaint was that the average score was only 72 points out of a possible 100.

What was odd about this reaction was that the average numerical score on the exam had absolutely no effect on the distribution of grades. The norm at the school was to use a grading curve in which the average grade was a B or B+, and only a tiny number of students received grades below a C. I had anticipated the possibility that a low average numerical score might cause some confusion on this front, so I had reported how the numerical scores would be translated into actual grades in the class. Anything over 80 would get an A or A-, scores above 65 would get some kind of B, and only scores below 50 were in danger of getting a grade below C. The resulting distribution of grades was not different from normal, but this announcement had no apparent effect on the students' mood. They still hated my exam, and they were none too happy with me either. As a young professor worried about keeping my job, I was determined to do something about this, but I did not want to make my exams any easier. What to do?

Finally, an idea occurred to me. On the next exam, I made the total number of points available 137 instead of 100. This exam turned out to be slightly harder than the first, with students getting only 70% of the answers right, but the average numerical score was a cheery 96 points. The students were delighted! No one's actual grade was affected by this change, but everyone was happy. From that point on, whenever I was teaching this course, I always gave exams a point total of 137, a number I chose for two reasons. First, it produced an average score well into the 90s, with some students even getting scores above 100, generating a reaction approaching ecstasy. Second, because dividing one's score by 137 was not easy to do in one's head, most students did not seem to bother to convert their scores into percentages. Lest you think I was somehow deceiving the students, in subsequent years I included this statement, printed in bold type, in my course syllabus: "Exams will have a total of 137 points rather than the usual 100. This scoring system has no effect on the grade you get in the course, but it seems to make you happier." And indeed, after I made that change, I never got a complaint that my exams were too hard.

In the eyes of an economist, my students were "misbehaving." By that I mean that their behavior was inconsistent with the idealized model of behavior that is at the heart of what we call economic theory. To an economist, no one should be happier about a score of 96 out of 137 (70%) than 72 out of 100, but my students were. And by realizing this, I was able to set the kind of exam I wanted but still keep the students from grumbling.

For four decades, since my time as a graduate student, I have been preoccupied by these kinds of stories about the myriad ways in which people depart from the fictional creatures that populate economic models. It has never been my point to say that there is something wrong with people; we are all just human beings—*homo sapiens*. Rather, the problem is with the model being used by economists, a model that replaces *homo sapiens* with a fictional creature called *homo economicus*, which I like to call an Econ for short. Compared to this fictional world of Econs, Humans do a lot of misbehaving, and that means that economic models make a lot of bad predictions, predictions that can have much more serious consequences than upsetting a group of students. Virtually no economists saw the financial

crisis of 2007–08 coming,* and worse, many thought that both the crash and its aftermath were things that simply could not happen.

Ironically, the existence of formal models based on this misconception of human behavior is what gives economics its reputation as the most powerful of the social sciences—powerful in two distinct ways. The first way is indisputable: of all the social scientists, economists carry the most sway when it comes to influencing public policy. In fact, they hold a virtual monopoly on giving policy advice. Until very recently, other social scientists were rarely invited to the table, and when they were invited, they were relegated to the equivalent of the kids' table at a family gathering.

The other way is that economics is also considered the most powerful of the social sciences in an intellectual sense. That power derives from the fact that economics has a unified, core theory from which nearly everything else follows. If you say the phrase "economic theory," people know what you mean. No other social science has a similar foundation. Rather, theories in other social sciences tend to be for special purposes—to explain what happens in a particular set of circumstances. In fact, economists often compare their field to physics; like physics, economics builds from a few core premises.

The core premise of economic theory is that people choose by optimizing. Of all the goods and services a family could buy, the family chooses the best one that it can afford. Furthermore, the beliefs upon which Econs make choices are assumed to be unbiased. That is, we choose on the basis of what economists call "rational expectations." If people starting new businesses on average believe that their chance of succeeding is 75%, then that should be a good estimate of the actual number that do succeed. Econs are not overconfident.

This premise of *constrained optimization*, that is, choosing the best from a limited budget, is combined with the other major workhorse of economic theory, that of *equilibrium*. In competitive markets where prices are free to move up and down, those prices fluctuate in such a way that supply equals demand. To simplify somewhat, we can say

* One economist who did warn us about the alarming rate of increase in housing prices was my fellow behavioral economist Robert Shiller.

that Optimization + Equilibrium = Economics. This is a powerful combination, nothing that other social sciences can match.

There is, however, a problem: the premises on which economic theory rests are flawed. First, the optimization problems that ordinary people confront are often too hard for them to solve, or even come close to solving. Even a trip to a decent-sized grocery store offers a shopper millions of combinations of items that are within the family's budget. Does the family really choose the best one? And, of course, we face many much harder problems than a trip to the store, such as choosing a career, mortgage, or spouse. Given the failure rates we observe in all of these domains, it would be hard to defend the view that all such choices are optimal.

Second, the beliefs upon which people make their choices are not unbiased. Overconfidence may not be in the economists' dictionary, but it is a well-established feature of human nature, and there are countless other biases that have been documented by psychologists.

Third, there are many factors that the optimization model leaves out, as my story about the 137-point exam illustrates. In a world of Econs, there is a long list of things that are supposedly irrelevant. No Econ would buy a particularly large portion of whatever will be served for dinner on Tuesday because he happens to be hungry when shopping on Sunday. Your hunger on Sunday should be irrelevant in choosing the size of your meal for Tuesday. An Econ would not finish that huge meal on Tuesday, even though he is no longer hungry, just because he had paid for it and hates waste. To an Econ, the price paid for some food item in the past is not relevant in making the decision about how much of it to eat now. An Econ would also not expect a gift on the day of the year in which she happened to get married, or be born. What possible difference can a date make? In fact, Econs would be perplexed by the entire idea of gifts. An Econ would know that cash is the best possible gift; it allows the recipient to buy whatever is optimal. But unless you are married to an economist, I don't advise giving cash on your next anniversary. Come to think of it, even if your spouse is an economist, this is probably not a great idea.

You know, and I know, that we do not live in a world of Econs. We live in a world of Humans. And since most economists are also human, they also know that they do not live in a world of Econs.

Adam Smith, the father of modern economic thinking, explicitly acknowledged this fact. Before writing his magnum opus, *The Wealth of Nations*, he wrote another book devoted to the topic of human "passions," a word that does not appear in any economics textbook. Econs do not have passions; they are cold-blooded optimizers. Think of Mr. Spock in *Star Trek*.

Nevertheless, this model of economic behavior based on a population consisting only of Econs has flourished, raising economics to that pinnacle of influence on which it now rests. Critiques over the years have been brushed aside with a gauntlet of poor excuses and implausible alternative explanations of embarrassing empirical evidence. But one by one these critiques have been answered by a series of studies that have progressively raised the stakes. It is easy to dismiss a story about the grading of an exam. It is harder to dismiss studies that document poor choices in large-stakes domains such as saving for retirement, choosing a mortgage, or investing in the stock market. And it is impossible to dismiss the series of booms, bubbles, and crashes we have observed in financial markets beginning on October 19, 1987, a day when stock prices fell more than 20% all around the world in the absence of any substantive bad news. This was followed by a bubble and crash in technology stocks that quickly turned into a bubble in housing prices, which in turn, when popped, caused a global financial crisis.

It is time to stop making excuses. We need an enriched approach to doing economic research, one that acknowledges the existence and relevance of Humans. The good news is that we do not need to throw away everything we know about how economies and markets work. Theories based on the assumption that everyone is an Econ should not be discarded. They remain useful as starting points for more realistic models. And in some special circumstances, such as when the problems people have to solve are easy or when the actors in the economy have the relevant highly specialized skills, then models of Econs may provide a good approximation of what happens in the real world. But as we will see, those situations are the exception rather than the rule.

Moreover, much of what economists do is to collect and analyze data about how markets work, work that is largely done with great care

and statistical expertise, and importantly, most of this research does not depend on the assumption that people optimize. Two research tools that have emerged over the past twenty-five years have greatly expanded economists' repertoire for learning about the world. The first is the use of randomized control trial experiments, long used in other scientific fields such as medicine. The typical study investigates what happens when some people receive some "treatment" of interest. The second approach is to use either naturally occurring experiments (such as when some people are enrolled in a program and others are not) or clever econometrics techniques that manage to detect the impact of treatments even though no one deliberately designed the situation for that purpose. These new tools have spawned studies on a wide variety of important questions for society. The treatments studied have included getting more education, being taught in a smaller class or by a better teacher, being given management consulting services, being given help to find a job, being sentenced to jail, moving to a lower-poverty neighborhood, receiving health insurance from Medicaid, and so forth. These studies show that one can learn a lot about the world without imposing optimizing models, and in some cases provide credible evidence against which to test such models and see if they match actual human responses.

For much of economic theory, the assumption that all the agents are optimizing is not a critical one, even if the people under study are not experts. For example, the prediction that farmers use more fertilizer if the price falls is safe enough, even if many farmers are slow to change their practices in response to market conditions. The prediction is safe because it is imprecise: all that is predicted is the direction of the effect. This is equivalent to a prediction that when apples fall off the tree, they fall down rather than up. The prediction is right as far as it goes, but it is not exactly the law of gravity.

Economists get in trouble when they make a highly specific prediction that depends explicitly on everyone being economically sophisticated. Let's go back to the farming example. Say scientists learn that farmers would be better off using more or less fertilizer than has been the tradition. If everyone can be assumed to get things right as long as they have the proper information, then there is no appropriate policy prescription other than making this information freely

available. Publish the findings, make them readily available to farmers, and let the magic of markets take care of the rest.

Unless all farmers are Econs, this is bad advice. Perhaps multinational food companies will be quick to adopt the latest research findings, but what about the behavior of peasant farmers in India or Africa?

Similarly, if you believe that everyone will save just the right amount for retirement, as any Econ would do, and you conclude from this analysis that there is no reason to try to help people save (say, by creating pension plans), then you are passing up the chance to make a lot of people better off. And, if you believe that financial bubbles are theoretically impossible, and you are a central banker, then you can make serious mistakes—as Alan Greenspan, to his credit, has admitted happened to him.

We don't have to stop inventing abstract models that describe the behavior of imaginary Econs. We do, however, have to stop assuming that those models are accurate descriptions of behavior, and stop basing policy decisions on such flawed analyses. And we have to start paying attention to those *supposedly irrelevant factors*, what I will call SIFs for short.

It is difficult to change people's minds about what they eat for breakfast, let alone problems that they have worked on all their lives. For years, many economists strongly resisted the call to base their models on more accurate characterizations of human behavior. But thanks to an influx of creative young economists who have been willing to take some risks and break with the traditional ways of doing economics, the dream of an enriched version of economic theory is being realized. The field has become known as "behavioral economics." It is not a different discipline: it is still economics, but it is economics done with strong injections of good psychology and other social sciences.

The primary reason for adding Humans to economic theories is to improve the accuracy of the predictions made with those theories. But there is another benefit that comes with including real people in the mix. Behavioral economics is more interesting and more fun than regular economics. It is the un-dismal science.

Behavioral economics is now a growing branch of economics, and

its practitioners can be found in most of the best universities around the world. And recently, behavioral economists and behavioral scientists more generally are becoming a small part of the policy-making establishment. In 2010 the government of the United Kingdom formed a Behavioural Insights Team, and now other countries around the world are joining the movement to create special teams with the mandate to incorporate the findings of other social sciences into the formulation of public policy. Businesses are catching on as well, realizing that a deeper understanding of human behavior is every bit as important to running a successful business as is an understanding of financial statements and operations management. After all, Humans run companies, and their employees and customers are also Humans.

This book is the story of how this happened, at least as I have seen it. Although I did not do all the research—as you know, I am too lazy for that—I was around at the beginning and have been part of the movement that created this field. Following Amos's dictum, there will be many stories to come, but my main goals are tell the tale of how it all happened, and to explain some of the things we learned along the way. Not surprisingly, there have been numerous squabbles with traditionalists who defended the usual way of doing economics. Those squabbles were not always fun at the time, but like a bad travel experience, they make for good stories after the fact, and the necessity of fighting those battles has made the field stronger.

Like any story, this one does not follow a straight-line progression with one idea leading naturally to another. Many ideas were percolating at different times and at different speeds. As a result, the organizational structure of the book is both chronological and topical. Here is a brief preview. We start at the beginning, back when I was a graduate student and was collecting a list of examples of odd behaviors that did not seem to fit the models I was learning in class. The first section of the book is devoted to those early years in the wilderness, and describes some of the challenges that were thrown down by the many who questioned the value of this enterprise. We then turn to a series of topics that occupied most of my attention for the first fifteen years of my research career: mental accounting, self-control, fairness, and finance. My objective is to explain what my colleagues and I learned along the way, so that you can use those insights your-

self to improve your understanding of your fellow Humans. But there may also be useful lessons about how to try to change the way people think about things, especially when they have a lot invested in maintaining the status quo. Later, we turn to more recent research endeavors, from the behavior of New York City taxi drivers, to the drafting of players into the National Football League, to the behavior of participants on high-stakes game shows. At the end we arrive in London, at Number 10 Downing Street, where a new set of exciting challenges and opportunities is emerging.

My only advice for reading the book is stop reading when it is no longer fun. To do otherwise, well, that would be just misbehaving.