There Is More to Behavioral Economics Than Biases and Fallacies

By Koen Smets
July 24, 2018

Compared to just a few years ago, the term *behavioral economics* has gained tremendous currency. To say it is on everyone's lips would be only a minor exaggeration. For the scientists and practitioners in the field, this emergence from relative obscurity should, at first glance at least, be a source of happiness.

One reason for this growing interest is the way behavioral economics has been presented and interpreted. Behavioral economics is, it seems, the field that confronts us with our deeply irrational selves. We are bamboozled by biases, fooled by fallacies, entrapped by errors, hoodwinked by heuristics, deluded by illusions.

(I am using *behavioral economics* as a catchall label. Behavioral and cognitive science has a wider scope, but much of it is concerned with how we make choices and decisions. That is both where the overlap with economics occurs and where its visible relevance to everyday life materializes.)

This is all very exciting, of course, and as a result we are knee-deep in articles and infographics that gleefully point out how flawed we really are. But is that really all there is to behavioral economics?

Behavioral economics emerged as a subfield of economics. Economists were discovering that people did not quite act and react like the consistently rational, self-interested, utility-maximizing agents in their neoclassical models. The deviations from the standard model had to be captured somehow, and psychology provided a basis for doing so. It is these deviations—or biases—that get the popular attention. That is, at best, a mixed blessing. To a worrying extent, biases have become the defining feature of behavioral economics.

This focus on biases is unhelpful in several ways. It fails to acknowledge that biases are broad tendencies, rather than fixed traits, and it oversimplifies the complexity of human behavior into an incoherent list of flaws. This leads to misguided applications of behavioral science that have little or no effect, or which can backfire spectacularly. We need to appreciate better the role biases play on the wider behavioral-economics stage.

The Troubles With Biases

A widespread misconception is that biases explain or even produce behavior. They don't—they *describe* behavior. The endowment effect does not *cause* people to demand more for a mug they received than a mug-less counterpart is prepared to pay for one. It is not *because of* the sunk cost fallacy that we hang on to a course of action we've invested a lot in already. Biases, fallacies, and so on are no more than labels for a particular type of observed behavior, often in a peculiar context, that contradicts traditional economics' simplified view of behavior.

A widespread misconception is that biases explain or even produce behavior. They don't—they *describe* behavior.

The conversation around biases is almost uniformly negative: they screw up our decision making, or undermine our health, wealth, and happiness. However, biases evolved with us, and for good reasons. When resources are scarce—as they were for tens of thousands of years in our existence as a species—loss aversion would have been a good bias to have. For early humans, losing a week's supply of food would have been much more significant than gaining an extra week's. Evolution can explain and indeed justify most of the biases that appear to make us so irrational today. And even today, loss aversion may not be quite so irrational.

Biases are not natural laws. They are broad tendencies, which are not uniformly shared by everyone. We are not all equally likely to respond to, say, social proof: some of us tend to be more conformist; others are more the rebellious kind. And even with these personal tendencies, the context can play a major role in how strongly a message will influence our behavior.

"A New Bias Every Day"

What are the consequences of this near obsession with biases? For one, the proliferation of biases masks the truth that human behavior is fluid and fuzzy. The use of discrete, distinct labels implies a rigor that is simply not there. There is no common guidance on how wide or narrow a bias can or should be defined. Despite valiant efforts like this one, there is no robust hierarchy, no underlying framework to indicate how biases relate to each other or to a coherent overarching theory. When research unearths an effect that doesn't quite fit any of commonly quoted biases, the simplest thing to do is to add another one to the pile.

This is not just a problem of presentation in the popular media. As Fabrizio Ghisellini, co-author of *Behavioral Economics – Moving Forward*, says, we have a discipline plagued with confusing definitions, unanswered questions and conceptual gaps—"a new bias every day."

But the propagation of an oversimplified understanding of human behavior is not the only issue. Starting a few years ago, an ongoing sequence of so-called failed replications has contributed to a climate of skepticism and erosion of trust in scientific research. Reports of yet another effect that does not replicate are met by a sensationalism akin to the hype that met the effect in the first place.

In some cases, follow-up studies have unmasked outright frauds, like Diederik Stapel, who fabricated and manipulated data to show that meat eaters are more selfish than vegetarians. Similarly, follow-up studies have exposed the dubious-but-perhaps-not-quite-fraudulent methods of people like Brian Wansink, whose papers have repeatedly been found to contain inaccuracies and errors undermining the validity of his conclusions.

But other "failed replications" are not quite so unequivocal once one looks behind the headlines. Human behavior is complex, and sensitive to context and circumstances. It is unwise to take an observed effect as gospel (we'll come back to that in a moment), but it is equally unwise to take a single failed replication as proof that the originally observed effect does not exist.

One example is the so-called priming effect. As long ago as 1996, John Bargh and colleagues found that when participants were presented with words associated with elderly people (bingo, retirement, etc.), they subsequently walked much more slowly than otherwise when they left the room. But how well does *this precise* "elderly–slow" effect replicate?

The proliferation of biases masks the truth that human behavior is fluid and fuzzy. The use of discrete, distinct labels implies a rigor that is simply not there.

A 2012 study by Doyen et al. *did* find a similar effect, but "only when experimenters believed participants would indeed walk slower." Bargh himself points at two successful replications. A first one, by Hull et al., found an interaction with self-consciousness; the second one, by Cesario et al., found an interaction with attitude towards the elderly. Statistician Andrew Gelman is skeptical, and suspects what he calls "noise mining": "There are so many ways to get 'p less than .05.'" But does that therefore mean that all *priming* effects, in all contexts, are imaginary?

Perhaps not. A 2016 meta-analysis by Weingarten et al. looked at 133 studies in which word primes were incidentally presented to participants and found small but significant priming.

Another example is that of the so-called *stereotype threat*—a phenomenon first described more than 20 years ago by Claude Steele and Joshua Aronson. The idea is that when a particular identity trait (race, gender, etc.) of a subject is emphasized, a stereotypical notion (e.g., physical or mental ability) is reinforced in the subject. Research by Steven Spencer, Steele, and Diane Quinn found "strong and consistent support" for the reasoning that stereotype threat can interfere with women's math performance (and that removing the threat can improve their results). For example, female students who had specify their gender prior to a math exam performed significantly worse than those who had to provide this information afterwards.

This particular instance, too, seems tricky to replicate. In 2004, researchers conducted an evaluation of students' performance on a range of mental ability tests in a field condition. One of their hypotheses was that asking students about ethnicity and gender would depress the performance of women on the quantitative tests. The

authors concluded that asking about ethnicity and gender did not affect performance, thus challenging the foundation of stereotype threat in this context.

In 2008, a different team revisited the data. It found that women benefited significantly when demographics questions were asked after the test, compared to before. Six years later, Paulette Flore and Jelte Wicherts carried out a meta-analysis on 47 comparisons of girls with and without stereotype threat and found a statistically reliable but small effect (d = 0.22), which, according to the authors, may be inflated as a result of publication bias.

In 2017, Bettina Casad, Patricia Hale, and Faye Wachs carried out a stereotype-threat experiment with 498 middle school girls. Their results paint a complex picture, not only with interactions with the course level (honors versus standard) but also with, for example, math attitudes and disengagement. They conclude that "stereotype threat is a real effect that occurs outside the laboratory."

There are a couple of lessons that we can draw from this. First, we should be circumspect when a study claims to have found a particular effect. A single study is interesting, but even the academic stamp of approval is not a guarantee for unconditional and universal applicability. Second, we should be mindful of the context dependence of much of our behavior. Claude Steele (who is preparing a preregistered replication of the stereotype threat effect on women's math performance) points out the importance of the situation in the role stereotype threat can play. This may be quite different for sixth-grade girls in Poland than for women doing elite coursework at a U.S. college. Even studies that aim to control for confounding variables will struggle to control for *everything* that could possibly interfere with the result.

Absolutist thinking is all too common in both directions: the effect applies universally and unqualifiedly, or—when it fails to replicate—simply never existed at all. This leads to unhelpful polarization and even animosity between positions. Yet the fact that our behavior seems so easy to influence should make us hesitate before assuming either that an effect observed in one set of circumstances will automatically apply in another context, or that a failure to replicate an effect automatically invalidates it. If there is real substance to an observed effect, failed replications help disentangle the complexity, improving clarity and precision.

A Little Knowledge

Aside from the proliferation and replication problems, a further consequence of overemphasizing biases, no less concerning, is starting to emerge: ill-informed (and failing) interventions. A few months ago, United Airlines hit the headlines, not, as most airlines do every so often, because it had screwed up a traveler's experience. Rather, United introduced an alternative incentive scheme for employees. Instead of the usual arrangements, in which every member of staff received a modest but not insignificant bonus, UA introduced a lottery. A few lucky employees would now get a large amount of cash, a luxury car, or vacation, while the others wouldn't get anything.

A little knowledge is, however, a dangerous thing. There is indeed a reason why a chance of winning something of a larger value might be a stronger incentive than the certainty of something of lesser value: we tend to overweigh small probabilities. A 1/100 chance of winning $100 appears more attractive than a certain $1.

Lotteries have shown positive results in health contexts, but they don't necessarily outperform fixed incentives. Harsha Thurumurthy et al. found that a guaranteed compensation of $12.50 worked significantly better as an incentive for medical circumcision (in a bid to control HIV) than did a lottery with a comparable expected value. A study by Aditya Vashistha et al. found that while participants said they preferred a lottery over a fixed incentive, the fixed incentive was far more effective.

This shows it's risky to assume tendencies (like the overweighing of small probabilities) apply unconditionally, or to overgeneralize findings from one context to another one, as indeed United Airlines found. Staff and unions were up in arms, and within days management had to revert to the original scheme.

It's risky to assume tendencies (like the overweighing of small probabilities) apply unconditionally, or to overgeneralize findings from one context to another one, as indeed United Airlines found.

Even when an experiment based on a behavioral economics insight is properly prepared and conducted, it may fail to deliver the hypothesized result. René de Wijk and colleagues [investigated](#) whether more accessible placement of wholemeal bread in a supermarket in the Dutch town of Veenendaal would increase its sales relative to white bread. This idea built on the much-cited study, reported in *Nudge*, in which the choice of fruit as a dessert increases when it is placed more accessibly on a cafeteria counter. However, the placement of the bread in Veenendaal appeared to have no influence on the customers' choices.

The Dutch experiment shows the right way to approach things, [even with effects that are widely accepted](#) to have general application: formulate a hypothesis and then *experiment*. But United Airlines' embarrassing experience illustrates the dangers of a quick fix. The simplification of scientific results—and I am explicitly also referring to the popular infographics gleefully displaying one-dimensional caricatures of biases and fallacies—has two potentially pernicious effects. First, they amplify findings that resonate with prior intuitions, making us believe they reflect profound and universal behavioral laws. The nuance quickly gets lost.

Second, because emphasizing the counterintuitive nature of certain behaviors is a surefire way to gain popular press, these infographics can go viral and do more harm than good. Surprise can make people renounce their prior beliefs and uncritically adopt the opposite view. A little knowledge is a dangerous thing, and much of the popular treatment of behavioral economics is, really, barely more than a little knowledge.

Armed with a sparkling new vocabulary of cognitive and behavioral effects, it's easy to see examples of biases all around us, and we fool ourselves into believing that we have become experts. We risk falling prey to confirmation bias. The outcomes of experiments appear obvious to us because we overlook the intricate nature of the full picture (or fail to notice unsuccessful replications). By simplifying human behavior into a collection of easily identified, neatly separate irrationalities, we strengthen our misguided self-perception of expertise.

Hubris in place of humility makes experiments superfluous—*science has said this is what happens*. And before you know it, your staff is on strike or your customers walk away.

It's Complicated

What does this all mean for behavioral economics? Maybe it's time to embrace the dreaded phrase, "It's complicated."

Oversimplification and overgeneralization obscure the actual complexity of human behavior. It's rare for a particular behavior to be a pure example of one specific cognitive effect. Most of the time we are subject not just to a multitude of contextual influences but also to multiple simultaneous effects that can combine in ways that are not immediately obvious.

Some of the recognized biases are opposites that can contradict each other. When we make a choice, are we influenced by what we saw *first* (priming or anchoring) or what we saw *last* (recency)? Are we influenced by what we know and are familiar with (status quo) or by what is new, shiny, and different (novelty)?

Even when there is no contradiction and effects can combine, it's still far from obvious how they will do so. Will people pursue a new initiative with great enthusiasm and perhaps obliviousness of the potential downsides (optimism), will they be held back by the perceived downsides (risk and loss aversion), or will they follow what their colleagues do (social proof)?

Behavioral economics is not magic: it's rare for a single, simple nudge to have the full desired effect.

Andrew Gelman [talks](#) of a "button-pushing" model: if you do X, then you can expect to see Y *every time*. He counters it with what he calls the "piranha argument": "There can be some large and predictable effects on behavior, but not a lot, because, if there were, then these different effects would interfere with each other, and as a result it would be hard to see any consistent effects of anything in observational data. The analogy is to a fish tank full of piranhas: it won't take long before they eat each other."

Perhaps even more importantly, amid the forest of biases and fallacies, we may forget that we still largely act and decide rationally. We often *[do respond to incentives](#)*, and we *do* carry out cost-benefit analyses—explicitly or implicitly.

From Complexity to (More) Clarity

Simplistic perspectives collapse complexity. A nice illustration of this is the so-called paradox of choice. In an often-cited experiment carried out in 2000 by Sheena Iyengar and Mark Lepper, offering shoppers an assortment of 24 types of jam led to fewer sales than presenting just six types. A rational agent would, of course, prefer *more* choice, so clearly people are irrational.

But as Sarah Whitley et al. found in a recent study examining the preferred size of choice sets in different situations, it's not that simple. People prefer fewer choices for utilitarian purchases and more choices for hedonic purchases. When we buy something only for its functional utility, we don't want to spend much time comparing various options—whatever does the job is good enough. When we are looking for something that will give us pleasure, in contrast, our preferences are more specific and pronounced, and this makes us more demanding.

A little reflection suggests that none of this is particularly irrational. There is a cost to choosing between different options. In the utilitarian setting the benefit is modest, hence we're not willing to incur a large cost. In the hedonic situation, the elevated benefit motivates us to spend more effort on making the best, and not just a good enough, choice.

Proceed With Caution

Raising general awareness of behavioral economics should be a good thing. It's good to be aware of new insights into the intricacy of our behavior, and of the ways in which it can help solve problems and pursue opportunities.

But there is a great need to beware of oversimplification. Learning the names of musical notes and of the various signs on a staff doesn't mean you're capable of composing a symphony. Likewise, learning a concise definition of a selection of cognitive effects, or having a diagram that lists them on your wall, does not magically give you the ability to analyze and diagnose a particular behavioral issue or to formulate and implement an effective intervention.

Behavioral economics is not magic: it's rare for a single, simple nudge to have the full desired effect. And being able to recite the definitions of cognitive effects does not magically turn a person into a competent behavioral practitioner either. When it comes to understanding and influencing human behavior, there is no substitute for experience and deep knowledge. Nor, perhaps even more importantly, is there a substitute for intellectual rigor, humility, and a healthy appreciation of complexity and nuance.

However, the plea for prudence is not just aimed at those of us who are new to the field of behavioral economics and who enthusiastically want to embrace its potential before having the necessary skills and experience. It is also aimed at behavioral scientists and practitioners. We need to continue to replicate, in different circumstances, and build up the model of human cognition and behavior—much like how a sculpture is made. Start with a rough wire mesh, add some dollops of clay, stand back and observe, trim and refine. Doesn't quite fit? Remove and start again. The difference is that a sculpture usually gets finished in the end. Behavioral economics will be a work in progress for a long time to come. But with the right sculptors, there is hope yet for a masterpiece.

Koen Smets

Koen Smets is an organization development adviser. He has been applying insights from behavioral economics in his work for many years, and writes frequently about human behavior in all its aspects. Koen holds a Master's degree in Engineering from Brussels Free University.Twitter