17
Reflections on Decision Theory: Response to Gilboa and March

David Schmidtz

17.1 What Do We Want from Decision Science?

Itzhak Gilboa observes (Chapter 16 of this volume) that the Bayesian approach holds that all uncertainty can and should be quantified by probabilities.[1] As I understand decision science, Gilboa is correct, and the point of the Bayesian postulate is to render decision theory apt for cranking out theorems. Theoretical tractability comes at a cost, however, in terms of practical realism.

Assigning probabilities, as if probabilities are known, is plausible enough, Gilboa would agree, in games of chance where information on probabilities and monetarily quantified utilities is readily available. It is less plausible to treat unknowns as equiprobable, however, when we can only guess at possible outcomes, when we can only guess at probabilities of such outcomes as we have the imagination to anticipate, and when we have hardly any clue how to gauge their utilities. Indeed, normal people typically do not even dream of trying to make decisions by using a formula to calculate maximum expected utility. We don't do that because we are rational. We know better than to handle real life's sometimes radical uncertainty by plugging pure guesswork into a toy formula.

So, on the one hand, I do not think being moral is closely connected to being a maximizer, as maximizing came to be construed in the twentieth century. On the other hand, I think that morality's departures from maximization do not cause any particular rift between morality and rationality because I also think being *prudent* generally has little to do with being a maximizer. I think then the founding philosophical challenge (as old as Plato's *Republic*) of identifying genuine, substantive connections between being rational and being moral requires working with *humanly* rational choice and humanly moral agency as they really are, not with mathematically tractable idealizations of them.

Here is a further and more positive thought on the prospects and practical limitations of decision science. I have never accepted the instrumentalist idea that the scope of rational choice is the selection of means to given ends, leaving us with nothing substantive to say about the rationality of ends themselves. To me, some ends obviously are better than others. Some ends make for better lives. But common-sense intuition aside, is there a non-question-begging *proof* that ends can be rationally chosen? Could instrumental rationality *all by itself*

(conjoined to descriptive truisms about human psychology) underwrite the choice of some ends rather than others, or some ends rather than none? In my earlier work, I did not argue that ends as such were rational, because that was never the question. The question is about ends as *chosen*. So, the real question is not whether final ends can be rational so much as whether they can be rationally *chosen*.

One of the exceedingly tricky aspects of the task of rationally choosing ends is that humanly rational choice is choice for beings whose preferences are a moving and evolving target. This is where the philosophical action is, but looking at rationality that way is not apt for cranking out theorems. As James March might have agreed (see his Chapter 15 in this volume), we tend to grow into the choices we make. We are outcomes as well as makers of our choices. That means that *whatever we think is at stake when we make choices, there is more at stake than that.* We tend to become people for whom our major life choices make sense. Sometimes this is a good thing, but it behooves us to make major life choices with one eye on the question of whether we want to give ourselves reason to become *that* kind of person.[2]

17.2 From Theory of the Good to Theory of the Better

In moral philosophy, we do something called "theory of the good." Curiously, there is no such topic as "theory of the better." Yet, consider how abstract the concept of "good" is and consider that we are not trained to notice how much more concrete the concept of "better" is. "Better" is *located*.

Our lack of a theory of the better has, I suspect, held back progress on what we call ideal theory. Prominent scholars such as A. John Simmons, not to mention Rawls, write as if, contra Amartya Sen, we obviously cannot understand *better* except by reference to *perfect*.[3] But not so. We rarely have any idea what to regard as perfect. By contrast, almost every decision we make is implicitly a decision about what to count as better. "Better" usually is defined by comparison to a starting point, not an endpoint. "Better" may involve imagining where we *want* to end up, and that will be a future-oriented idea. Yet, even that future thought will not be a conception of something perfect. It will itself be a conception of what to count as worthy of aspiration starting from here.

Perfection, conceived in any substantive way, is itself a moving target. Adult ideals evolve over the course of an examined life. Ideals can grow up. So, although of course we have targets, our targets are not *fixed* in a permanent way. We do not *need* such fixed targets. Thinking we can treat ideals as fixed targets and thinking that without fixed targets we could not make decisions is a way of thinking about ideals that fails to track what we know from experience.

---

[2] My first published thoughts on the respective arts of choosing strategies and choosing ends are in, respectively, "Rationality within Reason" (1992) and "Choosing Ends" (1994).

[3] Against Sen, Simmons says that, "We can hardly claim to know whether we are on the path to the ideal of justice until we can specify in what that ideal consists." Sen notes that we do not need to know that we are climbing the highest peak to know that we are climbing. Turning Sen's peak metaphor against him, Simmons concludes that, "which of the two smaller peaks of justice is the higher (or more just) is a judgment that matters conclusively only if they are both on equally feasible paths to the highest peak of perfect justice." See Simmons (2010: 34–35). See also Sen (2009). My reflections on this debate can be found in my "Nonideal Theory: What It Is and What It Needs to Be" (2011).

Decision science aims to crank out theorems. We have simple models of rational agency that are close enough to being recognizably human to have interesting implications. Simplification is one of the objectives of theorizing in general, including scientific theorizing. The kind of theory that illuminates is the kind that sticks its neck out and tries to boil the complexity of the human condition down to something simple. Simplifying—trying to understand the simple essence of things—does not really illuminate, however, unless we keep in mind that it exposes us to a certain contingency. Specifically, there is no simplification without risk of over-simplification, when the details we set aside as peripheral turn out to be central in a particular case.

When people spin off and begin to discuss messy complexity, it is not because they don't have the technical skill to crank out theorems. It's because they are trying to transcend the game of clever technicality and say something about the human condition that actually resonates.

17.3 When More Is Better

I noted that "better" is a concrete, located, and comparative idea. What about *more*? Should we assume "more is better"? In fact, "more is better" is a fruitful assumption notwithstanding the risk of it turning out to be an over-simplification in any given case. James March is right that assumptions such as this have put decision science and economics in a position to yield one insight after another into the human condition. Note, however, that "more is better" is not fruitful by virtue of being necessarily true. It is not necessarily true. Neither is "more is better" fruitful because it defines away our need to worry about when it's time to stop. Instead, what makes "more is better" useful is that it serves to remind us to think about when to stop.

It is a personal ideal to stop, think, and genuinely decide. Taking personal responsibility for deciding when we need to stop craving "more" generally will be an art that goes beyond the reach of decision science (including all the "nudging" in which people put such unwarranted faith nowadays). Computations may be part of a rational response to questions about when to stop, but another part of the response will be an existentialist responsibility to decide what to regard as worth wanting.

On one hand, there is a wide swath of human experience in which there is a grain of truth to the idea that more is better. On the other, when we test the assumption that more is better, and find cases that disconfirm, then take the hint that we need to reflect on what life is teaching us about when to stop, those are cases where we learn. Acknowledging grains of contingent truth and getting past tautologies frees us to draw contrasts. It sharpens our intuitive awareness of lines that *should not be crossed* on pain of ruining ourselves by wanting too much.

Obviously, when we assume for argument's sake that more is better, we should always ask, more of what? Strikingly, Adam Smith (1991 [1776]) could see what a terrible thing it can be to want too little. Smith understood that simply wanting "more" can diverge from wanting "better." That is to say, life is full of situations where unthinkingly wanting more could count as actually wanting too *little*. Smith saw workers laboring to bring services to

market that were good for their customers, but he also saw the dark side of that otherwise excellent social dynamic: namely, workers working overtime for trinkets and thereby failing to be good to themselves.

Smith saw market society as creating a possibility of leisure. Where a Robinson Crusoe living an altogether solitary life would not be able to make a pizza in 10,000 years, a person in favorable circumstances in any of today's commercial hubs can get a pizza for the price of a few minutes of labor. Smith was not, of course, saying citizens have become able to afford pizza, so much as that they have become able to afford leisure. We can afford to make time to think. We make a mistake when we take what could have been leisure time and spend that time commuting from a distant neighborhood where we can afford a larger house or working overtime for baubles. It is trite, yet still true, to say we treat wealth as our way of keeping score. Adam Smith could see that we do this, and he lamented it. Smith lamented that we don't spend our increasingly prodigious wealth stocking up on leisure. He would have agreed with the Society for Progress that the paradigm of business and management has focused overly on performance and not enough on progress.

However, Smith might also have told us that the proper place of business per se is not to actively *drive* a discussion about the value of leisure so much as simply to set the table for it. Commercial society has given us the time to constitute ourselves as a Society for Progress. Centuries ago, such a thing would hardly have been conceivable outside of a monastery. As Alan Krueger and colleagues said, money may be the currency of the economy, but time is the currency of life (Krueger et al., 2009). Commercial society has given us unprecedented latitude for choosing destinations of our own. It is a feature, not a bug, that commercial society cannot choose our destination for us. The mutual expectations that make us a society manage traffic. They help us to navigate. They help us to be advantageously predictable to each other; we learn how to stay out of each other's way. Above all, mutually understood expectations help us to cooperate, putting would-be cooperators in a position of knowing more or less who they can trust. Yet our ultimate destination—how we are going to use such opportunities and potentials as are actually available (however natural or socially constructed)—remains up to us.

17.4 Rationality and Bias

People can have different beliefs even when they have access to the same information. There are forms of confirmation bias that limit how good we can be as philosophers. The order in which we process bits of information affects what it takes for us to be ready to take the next bit at face value. Therefore, you could take two identical clones and give them identical sets of information but feed them the identical bits of information in a different order. By the lights of decision science, that should not make a difference in theory, but in fact it does. The first bit that comes in, if it is plausible, is provisionally accepted, but it has to cohere with the next bit. If it does not cohere with the next bit, and if the next bit is only equally plausible, but not so plausible that it over-rules the first bit, then the second bit is rejected and the first belief, the incumbent belief, is retained. But if the second clone gets the second bit first, then the second clone retains that bit instead, and the two clones end up with different beliefs precisely because they each processed the same information (in a different order) in real

time. The fact that human beings process information one bit at a time will result in the two clones ending up with different beliefs.

A human being is a process, unfolding in real time. Every judgment is temporally located, and our history of accumulated judgment adds up to our path-dependent framework for further judgment. We do not have the option of "all-things-considered" judgment. The fact that information processing is a *process* limits how good we can be at science or philosophy. We don't have to be biased to end up being skewed. All it takes is needing to process information in a temporal sequence. We can't do much about that.

Then it gets worse. A recent study of people's ability to process information about changing trends in global temperatures shows our processing ability being deranged by the most subtle evidence of an ideological orientation. A grayed-out stylized cartoon of an elephant or a donkey in the background of the graph of changing temperature threw subjects off, regardless of which party they identified with. To some, a background image says relax, give this information the benefit of the doubt; the same image leads others to tense up, assume bad motives, and seek reasons to discount the information (Guilbeault et al., 2018).

We process information as if information processing were a team sport. That is a problem. We need to learn to be exceedingly careful about classifying people, especially ourselves. Decisions we make about how to pigeonhole people compromise us as we go, undermining our ability to listen and learn. We imagine we are being responsible when we treat certain sources with skepticism, but often we are kidding ourselves. The uncomfortable truth is that people who see things differently are the ones from whom we have the most to learn.

17.5 Rationality as a Process

We can treat variables as variables. But we have a theoretical reason to treat variables as fixed. It makes our models look more like controlled experiments. But treating our *preferences* as exogenous is a simplifying assumption that can easily mislead.

Jim March observed that the process of pursuing an objective often transforms the values that justify it. Again, this is an idea that is not so amenable to the project of cranking out theorems, but we have over-rated the value of the theorem game. Being able to crank out theorems surely is worth something, but it is no substitute for the messy and humbling project of reflecting on the human condition in such a way that some of our isolated and individual bits of wisdom will rub off on each other and add up to something from which we can all take comfort and inspiration.

Humanly rational choice—rational choice for beings whose preferences are a moving and evolving target—is where the philosophical action is, but looking at rationality that way is not apt for cranking out theorems. We are outcomes as well as makers of our choices. That means that whatever we think is at stake when we make choices, there is more at stake than that.

We are all spending our lives writing a novel, a work of more or less historical fiction. We are characters in that novel. More than anything else, we want our story to be a story about it mattering that we were here. Look at any Facebook page. The iPhone camera is a stunning commercial success in part because it enables us to tell our life story in selfies.[4]

We might be heroes; that's up to us. There is an existentialist perspective, not easily ignored or discounted, from which it looks like nothing really matters at the end of the day. Yet, if

---

[4] I owe the point to Liz Anderson, in conversation, spring 2019.

anything matters to us personally, it is being able to be at peace with people coming after us, reading that novel, hearing our story, and knowing who we were.

Summarizing, and repeating: rationality and morality are not the same. No one should want to spend their lives proving that rationality and morality are tautologously co-extensive, rendered compatible by definition. In fact, bringing the two together is an achievement, not a given. The compatibility, indeed the synergy, between rationality and morality is a hard-won contingency, not a clever philosophical trick. For all that, however, the achievement is there for the taking. It is a robust and frequently realized possibility: an examined life can turn out to be a life worth living.

References

Guilbeault, Douglas, Joshua Becker, and Damon Centola. "Social Learning and Partisan Bias in the Interpretation of Climate Trends." *Proceedings of the National Academy of Sciences* 115 (2018): 9714–9719.

Krueger, Alan B., Daniel Kahneman, David Schkade, Norbert Schwarz, and Arthur A. Stone. "National Time Accounting: The Currency of Life." In *Measuring the Subjective Well-Being of Nations: National Accounts of Time Use and Well-Being*, edited by Alan B. Krueger, 9–86. Washington, DC: National Bureau of Economic Research, 2009.

Plato. *Republic.* Translated by G. M. A. Grube. Revised by C. D. C. Reeve. Indianapolis, IN: Hackett Publishing Company, 1992.

Schmidtz, David. "Rationality within Reason." *Journal of Philosophy* 89 (1992): 445–466.

Schmidtz, David. "Choosing Ends." *Ethics* 104 (1994): 226–251.

Schmidtz, David. "Nonideal Theory: What It Is and What It Needs to Be." *Ethics* 121 (2011): 772–796.

Sen, Amartya. *The Idea of Justice.* Cambridge, MA: Harvard University Press, 2009.

Simmons, A. John. "Ideal and Nonideal Theory." *Philosophy and Public Affairs* 38 (2010): 5–36.

Smith, Adam. *The Wealth of Nations.* London: David Campbell Publishers, 1991 [1776].