

Ockham's Razors: A User's Manual, by Elliott Sober. Cambridge: Cambridge University Press, 2015. Pp. x + 314.

Elliott Sober's first book, *Simplicity* (1975), defends the view that the simplicity of a theory or hypothesis is a measure of its informativeness – roughly, simpler theories require less new information to be added to them to answer relevant questions of interest. While this measure of simplicity is question-relative, it is still what you might call a *global* view of simplicity – simplicity means the same thing across different scientific problems and it is always an epistemic virtue. Ockham's Razor is just good scientific reasoning. Sober's 1988 book *Reconstructing the Past: Parsimony, Evolution, and Inference* argues against this global conception of simplicity and replaces it with a local one. Here, in one context parsimony (now roughly synonymous with simplicity) means one thing while in another context it may mean something different. Similarly, whether the more parsimonious hypothesis is to be preferred is also a local matter; it depends on various background assumptions in play, something which can differ between contexts. Ockham's Razor is *not* a universally valid inference rule. Sober's 1990 'Let's Razor Ockham's Razor' pushes this view even further by suggesting that perhaps there is no reason to continue using Ockham's Razor at all. The basic reason is that parsimony is not a *sui generis* epistemic virtue. When parsimony goes along with something of basic epistemic value, then the parsimonious theory is to be preferred. If it does not, then parsimony is of no value.

Sober's most recent book *Ockham's Razors: A User's Manual* brings together much of his previous work on parsimony, substantially expands his previous arguments and presents a host of new ones. It continues the localist and reductionist traditions of *Reconstructing the Past*. As Sober puts it, 'If parsimony contributes to the achievement of some more fundamental epistemic goal, I am all for it. If it does not, I am not' (p. 149). However, in this work, Sober argues that there is no need to 'razor the razor' as he previously suggested. Just because parsimony is not itself a fundamental theoretical virtue, it does not follow that it is not valuable to understand and study it: 'The epistemic relevance of parsimony does not require that parsimony be an end in itself' (p. 149). On this point, I believe that Sober is correct. Parsimony could be a relevant consideration even if it *always* reduces to something else. But it is still a relevant question to ask whether in fact it does always reduce to something else. Sober argues yes, it always does (or perhaps when it doesn't, then it is not of any value). Sober's reductionism entails that we *could* razor the razor if we wished without changing which inferences were good ones. I am not so sure that we could do so and I will return to this point near the end of this review.

As well as reinforcing some of his previous arguments, Sober frames his discussion within a larger historical and philosophical context. He also argues for a number of substantial theses which are important and interesting

regardless of how they relate to discussions about Ockham's Razor. It is an impressive work – well worth reading for philosophers and scientists from a number of different fields and essential reading for philosophers of science interested in scientific epistemology.

Let's begin with the basic question of just what Ockham's Razor actually is. According to Sober, 'Ockham's Razor, the principle of parsimony, says that a theory that postulates fewer entities, processes, or causes is better than a theory that postulates more, so long as the simpler theory is compatible with what we observe' (p. 2). Despite this seemingly clear definition, Sober goes on to say that 'different thinkers have meant different things by parsimony and different justifications for principles of parsimony have been constructed. My goal in this book is to describe this diversity and to determine when parsimony is relevant and when it is not' (p. 2).

By equating Ockham's Razor with a kind of quantitative and ontological parsimony in this way, Sober is immediately setting aside discussion of various aspects of simplicity, such as the beauty and elegance of a theory which might be called 'syntactic simplicity' (Sober 2001). I do not think it is a problem for Sober that he ignores these kinds of uses, but it is worth noting that his pluralistic conclusions about the varieties of ways in which the razor can be used could easily be broadened.

Having got clear about what the principle of parsimony actually says, it is time to ask whether it is justified. To cut to the chase, Sober says in his introduction that '[p]arsimony arguments that draw conclusions from the fact that one theory is more parsimonious than another differ from each other at two levels. First, some of them succeed while others fail. Second, the successful arguments succeed for different reasons, and the unsuccessful arguments go wrong in different ways' (p. 3). The book goes on to make a few general claims about Ockham's Razor, but each claim is justified by a detailed look at a variety of special cases examined in their local context.

Chapter 1 features a collection of short historical examples including Aristotle, William of Ockham, Copernicus, Descartes, Leibniz, Newton, Hume, Kant, Mill, Whewell, Maxwell, and Morgan. In looking at a variety of uses of Ockham's Razor, we can see right away some important differences between them. To help categorize the uses, Sober splits the principle of parsimony into two principles: the 'razor of silence' and the 'razor of denial'. The distinction is supposed to parallel the difference between agnosticism and atheism. The stated purpose of the chapter is not just to look at historical cases where a principle of parsimony has been used, but to look at the various attempts to justify such a principle. While the justifications differ across these cases, theological justifications are common, as is the idea that the principle is justified because nature itself is simple. Sober argues that neither of these traditional justifications is acceptable.

Chapter 2 moves the discussion to the twentieth century where discussions of simplicity take a 'probabilistic turn'. Sober argues that we can find two

distinct ‘parsimony paradigms’, which show that parsimony is epistemically relevant. One is the Bayesian paradigm on which parsimony can be justified when it lines up with likelihood – that is, the probability of the observations given the hypothesis. A second paradigm is frequentist. Here parsimony can play a role in model selection criteria such as the Akaike Information Criterion, where parsimony counts the number of adjustable parameters in a model. Since reducing the number of adjustable parameters in a model reduces overfitting, parsimony in this sense is relevant to the predictive accuracy of a model.

Understanding these paradigms requires understanding some technical details, so the chapter features a very short introduction to probability theory, Bayes’ Theorem, models and model selection theory. An example of the likelihood justification for parsimony is examined through the lens of Reichenbach’s common cause principle, where Sober considers the common cause hypothesis for two correlated events to be more parsimonious than the hypothesis of two separate causes. Sober then investigates the conditions under which the more parsimonious hypothesis matches the one with the higher likelihood. (The likelihood of a hypothesis H is $P(O|H)$ – the probability that the hypothesis confers on the data.) It turns out that they often (but not always) match.

Sober (2001) describes a third parsimony paradigm: sometimes parsimony is a stand-in for prior probability. This would also be a kind of Bayesian justification. In fact, it is sometimes treated as ‘the standard’ Bayesian justification for Ockham’s Razor (Earman 1992). While in one footnote Sober suggests that this paradigm ‘plays third fiddle’ (p. 141), in his more comprehensive discussion of the connection between simplicity and prior probability, Sober argues that this potential justification for Ockham’s Razor is fatally flawed. Siding with Popper, Sober argues that it is impossible for simpler hypotheses always to have higher probabilities than more complex ones for the reason that simpler hypotheses often entail more complex ones and so therefore the more complex ones cannot have lower probabilities (though he does not accept Popper’s claim that simpler hypotheses are more falsifiable).

Chapter 3 looks at the connection between parsimony and likelihood much more deeply through an investigation of the problem of phylogenetic inference and the inference methods of cladistic parsimony and maximum likelihood which are frequently used in such problems. *Reconstructing the Past* provided an in-depth look at this connection. Here Sober reiterates a number of previous points, including how to think about which background assumptions are sufficient or necessary for the two methods to favour the same phylogenetic trees. He also updates the reader by pointing to a number of new discoveries about this relationship that have been made in the intervening time.

Chapter 4 is a fascinating look at the scientific question of whether chimpanzees are mind-readers – that is, whether they form mental representations

of the mental states of others. To some authors, it seems that there is a kind of default presumption in favour of the view that they have no such abilities. That this is the default follows from Morgan's Canon: 'In no case may we interpret an action as the outcome of the exercise of a higher psychical faculty, if it can be interpreted as the outcome of the exercise of one which stands lower in the psychological scale' (Morgan 1894, p. 53). This is plausibly an instance of the principle of parsimony, though as Sober points out, Morgan actually thought that positing the higher faculties was the simpler theory but wanted to warn against it. This idea that one theory might be simpler than another on one set of grounds and less simple on some other set has important consequences.

Researchers have performed a number of different experiments which seem to indicate that chimpanzees can mind-read. What we observe are the environments in the experimental setup and the behaviours that the chimpanzees exhibit. This leaves open the question of the psychological mechanisms leading to the behaviours and so is an instance of what Sober calls 'black-box inference'. Povinelli and Vonk (2004) argue that there is a logical problem with positing an ability of the chimpanzees that we can't directly observe: it is possible that the chimpanzees are predicting the behaviour of the other chimpanzees directly based on the environmental stimuli they can see. Positing mind-reading is positing an extra intervening variable (namely, representations of mental states) between the environmental features and the chimpanzees' predictions of other chimpanzees' behaviours and, therefore, mind-reading is an unnecessarily complex hypothesis.

Tomasello and Call (2006) argue that positing that chimpanzees can mind-read allows us to explain the results of different experiments with a single unifying explanation. Unification and parsimony often go hand in hand and it is natural to think that on this basis, mind-reading is actually the simpler explanation. Another way to view this issue is to argue, as Whiten (1996) does, that the mind-reading hypothesis gives us a natural way to represent the causal story in a more economical way with fewer causal arrows. Sober ultimately rejects Whiten's arrow counting argument and argues instead that when we think about the two hypotheses as models, mind-reading has an extra adjustable parameter and so is less parsimonious. Of course, mind-reading is still more unifying and so Sober treats the case as one where parsimony and unification are in conflict. While I agree that parsimony and unification can conflict, it seems to me that it is very important to recognize that parsimony considerations can conflict with each other and not just considerations having to do with Ockham's Razor broadly understood. In other words, not only is it not clear what the epistemological impact is when one theory is simpler than another, it is often not clear which theory is simpler in the first place. This case seems a perfectly good example of this. If Tomasello and Call had argued that mind-reading hypothesis is better on the grounds of simplicity, rather than explicitly citing unification, it would make just as much sense.

This dispute over parsimony and unification has led to something of a stand-still in the debate over whether chimpanzees can mind-read. Rather than fighting over methodological principles, Sober proposes a new type of experiment that he hopes will shed some light on the debate.

In chapter 5, Sober discusses Ockham's Razor in connection with debates in philosophy, including debates about: the existence of God, the mind-body problem, the causal efficacy of the mental, moral realism, Plantinga's evolutionary argument against naturalism, nominalism about mathematics, and solipsism. Apropos the above discussion, when I saw this list in the contents I naturally assumed that solipsism was the parsimonious hypothesis. What could be more ontologically parsimonious than getting rid of the external world?! But Reichenbach himself argued that there was a common cause argument for the existence of the external world and harking back to the arguments of chapter 2, Sober treats this as a parsimony argument. As it turns out, Sober argues that this particular common cause argument fails since even if we assume there is a common cause, it may well be internal to our minds. In addition to any common cause considerations in favour of solipsism, like the chimpanzee mind-reading hypothesis, the external world hypothesis is more unifying, which makes it seem simpler in certain respects. The problem of the external world does seem precisely the kind of problem that is not settled on empirical grounds. This is exactly where Ockham's Razor is supposed to shine. But since there are multiple, competing parsimony considerations in this case, we have good reasons for thinking that Ockham's Razor will not solve this particular problem for us.

What about the existence of God? Sober focuses only on the problem of evil. As he sees it, atheism is more parsimonious than theism and the law of likelihood shows that the observations of evil in this world evidentially favour no PKG (all-powerful, all-knowing, and perfectly good) God over the existence of a PKG God. Thus, this is one of the justified uses of the principle of parsimony, since it falls under one of the parsimony paradigms.

Perhaps the most common argument against the existence of God is the problem of evil. But a close competitor is undoubtedly some version of Ockham's Razor. When asked 'Why not believe in God', atheists such as Hitchens (2007) sometimes recount the story (likely apocryphal) involving the great French physicist and mathematician Pierre-Simon Laplace. Laplace is asked by Napoleon why his book *Treatise on Celestial Mechanics* made no mention of God. He is supposed to have replied 'Je n'avais pas besoin de cette hypothèse-là'. ('I had no need of that hypothesis'.) The atheist's move from 'God is not needed' to 'God doesn't exist' is the razor of denial in action.

But like chimpanzee mind-reading and the external world, the hypothesis that there is a God has often been thought to be unifying and, in fact, a more simple hypothesis than atheism (Swinburne 2004). Once we start looking, it is easy to find cases where it is unclear which of two competing hypotheses is

actually simpler. Often this is because we tend to associate particular features with simplicity (the paucity of types of entities and unification, for example) and if we do this, it is easy to see that these can come into conflict. All we need to do is simply introduce a new entity which allows a unification between two previously distinct explanations. Of course if the entity introduced is clearly ad hoc, the ‘unifying’ hypothesis will not be a serious one. But if there is some reason to think it is plausible (as in the case of chimpanzee mind-reading or God) we have the recipe for a serious debate. These kinds of trade-off cases cannot be decided a priori; they must be looked at on a case-by-case basis. Sober’s detailed discussions of apparently unrelated cases is of great philosophical and scientific interest and ultimately justifies his view in the ‘localness’ of the value of Ockham’s Razor. But direct arguments like this one lead to the same conclusion.

Now let’s look at some of the more global conclusions in the book. As I mentioned earlier, Sober splits the principle of parsimony into two principles: the ‘razor of silence’ and the ‘razor of denial’. The distinction is supposed to parallel the difference between agnosticism and atheism. The exact statements of the principles can be found in Sober (2009, p. 128):

Razor of denial. If your evidence does not discriminate between ‘*X* exists’ and ‘*X* does not exist’, you should deny the former and affirm the latter.

Razor of silence. If your evidence does not discriminate between ‘*X* exists’ and ‘*X* does not exist’, you should suspend judgment about both.

Since they give incompatible recommendations in the same cases, the principles themselves are incompatible. When discussing whether some application of Ockham’s Razor is justified, Sober often points out whether it is an application of the razor of silence or of the razor of denial. So which, if either, of these principles is justified?

Sober suggests that the razor of silence has an obvious justification. For any probability function and any *X* and *Y*, $\Pr(Y \& X) \leq \Pr(Y)$. (I have made a few trivial notational changes for readability.) Therefore, ‘slicing away’ *X* by declining to assert or deny it will not decrease probability and will often increase it. As Sober puts it, ‘silence reduces your risk of error. The razor of silence has a simple Bayesian rationale’ (p. 71). I disagree.

First, *Y* & *X* and *Y* are not actually competing hypotheses. Why shouldn’t I believe both? Or neither? While Sober himself is not a Bayesian and makes no recommendations in the book about what we ought to believe, Sober implies that this basic fact about probability tells us that if you are a Bayesian, then rather than believing *Y* & *X* you should believe *Y* and should suspend judgment on *X*.

But what if *X* is actually true? Then by suspending judgment on *X* rather than believing it, you are believing fewer true things. Whether it is rational to believe *X* obviously depends on the evidence that you have for it. In some sense, suspending judgment on *any* of your beliefs at all would increase probability

(namely, by increasing the probability that the conjunction of all your beliefs is true). It certainly doesn't follow that there is a Bayesian rationale for doing so.

The razor of silence is a principle about rational belief (or rational acceptance). Thus, in order to provide a rationale for this razor (or the razor of denial), you need a theory of rational belief. Bayesianism is not such a theory. Now, it might seem as though there is an obvious connection between Bayesianism and rational belief: Bayesianism tells us that we should have a degree of belief in every proposition and then we can simply accept some version of the Lockean thesis which connects degrees of belief to beliefs. For example, a simple version of this might be 'It is rational to believe anything with probability $> .5$, rational to disbelieve anything with probability $< .5$ and rational to suspend judgment on propositions with probability $= .5$ '. I am not claiming that this is a good principle of rationality, merely that it is the *kind* of principle which could relate to the razors of silence and denial.

So, can such a principle justify the razor of silence? Of course, the key here is the antecedent: 'your evidence does not discriminate between "X exists" and "X does not exist"'. But what exactly does this mean? If it means that $\Pr(X \text{ exists}) = \Pr(X \text{ does not exist}) = .5$ then it does seem reasonable to suspend judgment on X and our Lockean principle would entail this. However, this has nothing to do with the probabilities of conjunctions and it seems odd to call this a 'Bayesian rationale' for the razor. In addition, equal probability as an interpretation of 'the evidence does not discriminate' is problematic, since Sober often rejects the use of prior probabilities as unjustified. Sober's favoured principle of evidence is The Law of Likelihood. It says that O favours H₁ over H₂ to the extent that $P(O | H_1) > P(O | H_2)$. The natural consequence is that 'evidence doesn't discriminate between X and $\sim X$ ' is $P(E | X) = P(E | \sim X)$. This equality is consistent with X having any probability between 0 and 1. For example, my evidence doesn't discriminate between this ticket winning this large, fair lottery versus any other ticket winning. But in such a lottery, it is perfectly reasonable for me to have a very low degree of belief that this ticket will win. It therefore seems unreasonable that I should suspend judgment on this proposition. Instead, I should believe that I will lose.

While this example doesn't immediately deal with the existence of entities, it is easy to construct parallel cases that are about ontological matters. For a Soberesque example, despite its high likelihood, I should not believe that there are gremlins bowling in my attic (Sober 2008). In fact, I think that it is appropriate that I believe there are not gremlins rather than merely suspending judgment. For a plausible example that moves from low to high probability, my degree of belief that there is currently at least one ant within 100 feet of the front door of my house is extremely high – though it isn't clear that my current evidence actually discriminates between the two relevant hypotheses. I have some inductive evidence about the past, plus a very general understanding of how common ants are where I live. But I have

no direct evidence at all about whether there is *right now* an ant in that location. After all, I am in my office away from my home and my belief is not sensitive in any way to whether or not there is an ant within 100 feet of the front door of my house.

Now, certainly, Bayesianism by itself doesn't recommend that I suspend judgment on whether there are gremlins in my attic or ants near my front door. Contra Sober, the razor of silence has no Bayesian justification.

What about the razor of denial? Unlike the razor of silence, Sober never suggests that the razor of denial has any kind of general justification. Since it is directly incompatible with the razor of silence, any justification for silence is thereby an argument against the razor of denial (and vice versa). In its full generality, the principle has serious problems. For example, if 'the evidence doesn't discriminate' means that $P(X) = P(\sim X) = .5$, it is certainly sometimes reasonable to suspend judgment rather than believe $\sim X$. If we take 'the evidence doesn't discriminate' to mean $P(O|X) = P(O|\sim X)$, then we have cases where $P(X)$ is quite high. But instead of thinking of the principle in full generality, we can instead consider special cases of it.

Consider completely superfluous hypotheses. One clear conclusion is that our best scientific theories should not include such entities. Sometimes a theory doesn't mention some particular entity and we take that merely to mean that the theory is silent on the question of whether the entity exists. But this is not always true. Consider the famous case of Einstein's theory of special relativity, which is more parsimonious than Lorentz's theory since special relativity denies the existence of the aether. Sober (1981) uses this example to argue that the razor of silence is incorrect since it conflicts with the razor of denial, which is sometimes justified!

In the case of God, as mentioned before, it is common for atheists to cite Ockham's Razor-like justifications. Why not believe in God? Sober considers the Problem of Evil – one common reason people cite for being an atheist. But here is one way out of that particular problem: 'God' is not actually perfectly good. Now I myself do not believe that there is an all-powerful, all-knowing being. (Set aside whether he is perfectly good.) What exactly justifies my belief that there is no such being? To bring out the puzzle clearly, imagine that such a being might not ever actually causally intervene in the world. On a purely likelihoodist understanding of evidential discrimination, it is hard to see how my evidence discriminates between this kind of God existing versus not existing. Perhaps the best thing to say is that the likelihoods are *inscrutable* rather than equal, but this does not give us a reason to deny this God's existence.

At this point, a reader might suggest that I don't need Ockham's Razor to justify my disbelief in such an entity. The prior probability of such a thing existing is so low that when combined with the fact that I have no evidence for or against it, it is rational to believe such a thing does not exist. I agree,

though Sober (2008) is sceptical that prior probabilities for the existence of beings like this can be justified. If we do accept that the prior probability should be low, the obvious next question is: *Why does the existence of such a being have a low probability in the first place?* I suggest that whatever the answer to that question is, it will be a pretty good summary of (some version of) Ockham's Razor.

In chapter 2, Sober points out that sometimes we call an explanation the 'simplest' when we mean that it is the best or most probable explanation. But where does this high probability come from? In standard cases, such as asking whether a patient has a rare or more common disease, simple induction based on past frequencies gives us rational justification for these priors. Sober makes the distinction between 'first priors' and 'non-first priors', which (I think) is meant to capture the idea that there is nothing necessarily unparsimonious about the hypotheses that have low prior probabilities if they are based on conditioning on past evidence. But what about first priors? This is precisely where Bayesians have thought that simplicity can matter. For example, Swinburne concludes his 2009 book *Simplicity as Evidence for Truth* with the following: 'To summarise the claims in a nutshell: either science is irrational (in the way it judges theories and predictions probable) or the principle of simplicity is a fundamental synthetic a priori truth' (p. 56). In some cases, I agree. I just disagree with Swinburne when he says that the God hypothesis is simple. I think it is much simpler to deny the existence of God and our respective beliefs about simplicity in part ground our prior probabilities for the existence of God.

If you don't like the God example, think of the mysterious creatures that I just made up: 'can't see ums'. If they did exist, I wouldn't expect to have any evidence for their existence. After all, they are scared of humans and incredibly good at hiding. It seems to me completely appropriate to believe that they don't exist rather than just suspending judgment. Evidence is a tricky notion, but I am inclined to say that I do not have any evidence that such creatures do not exist. Rather, there is a default presupposition of rational thinking that you should assume that totally ad hoc hypotheses with nothing to count in their favour can be summarily dismissed.

With cases such as these, we are getting closer to understanding why the razor of denial sometimes seems justified. First, following Barnes (2000), let's distinguish between a mere anti-quantity principle (positing fewer entities is better) and the weaker anti-superfluity principle, which says that theories should not contain superfluous components. (This is weaker since there are often theories with more components than another theory, but none of the components are superfluous.)

Add to this another common distinction in the literature: the distinction between qualitative and quantitative parsimony. Lewis (1973, p. 87) distinguishes between qualitative and quantitative parsimony where qualitative parsimony is a matter of keeping down the number of *kinds* of entities

while quantitative parsimony is a matter of minimising the number of entities within a kind. Lewis holds that while qualitative parsimony is a good thing in a philosophical or empirical hypothesis, quantitative parsimony is not relevant. There is a lively debate about the status of quantitative parsimony, but the authors all seem to agree with Lewis that there is a basic presumption in favour of qualitative parsimony (Baker 2016, Jansson and Tallant 2016, Nolan 1997). There is a clear pattern from the history of science: when a special type of entity is shown to be explanatorily superfluous, it is eliminated from scientific theories. There are numerous examples, such as privileged spatial and temporal locations, absolute space and time, phlogiston, the aether, vital forces, the soul and God. Many of these apparently led to specific predictions, which then turned out false. But at some point, all were posited in such a way that we don't have direct observational evidence that they do not exist. But when they are thought to be explanatorily impotent and superfluous it is claimed that they do not exist. When someone wants to dispute this conclusion, they argue that in fact these entities are not superfluous. Rather, they are explanatorily powerful and therefore we do have evidence for their existence. No one ever responds in the following way: 'You are correct, we have no evidence for vital forces. However, we should remain agnostic. Perhaps they exist, perhaps not'. When we do not know whether to believe in a special type of entity or not (for example, we aren't sure if substance dualism is true) it is because we are not sure whether or not non-physical substances are in fact superfluous.

These 'denial inferences' seem to me completely justified. Specifically, I would argue that conditionals such as 'If the soul is explanatorily superfluous, then it doesn't exist' are justified. This seems to me a clear use of Ockham's Razor. But here parsimony does not mirror likelihood or predictive accuracy. In chapter 5, Sober argues that we can use model selection techniques to see why the mind-body identity theory would be more predictively accurate than substance dualism: it uses fewer adjustable parameters. He might then suggest that similar considerations will work in the 'undetectable God' case and other cases of explanatorily superfluous entities. But even if he is right about dualism (and I don't think he is, though reasons of space prevent me from arguing about this here), model selection techniques and statistical theorems such as Akaike's theorem are not sufficiently broad to cover all the kinds of cases that the anti-superfluity principle will apply to.

Perhaps there is some other fundamental epistemic virtue that parsimony mirrors in these cases. Sober leaves open that there may well be other parsimony paradigms. But I don't know what they are. Sober is certainly correct that there are a wide variety of uses and justifications (good ones and bad ones) for Ockham's Razor. But Sober also concludes that there is no non-reductive justification for the principle in any case. If we split cases finely so that we are looking at qualitative parsimony and cases of superfluous entities,

I am not so sure. Here, perhaps Sober's pluralism should make room for one more parsimony paradigm – a non-reductive one.

References

- Baker, Alan 2016, 'Simplicity', *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2016/entries/simplicity/>>.
- Barnes, E. C. 2000, 'Ockham's Razor and the Anti-Superfluity Principle', *Erkenntnis* 53, pp. 353–374.
- Earman, John 1992, *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory* (Cambridge, MA: The MIT Press).
- Hitchens, Christopher 2007, *God is not Great: How Religion Poisons Everything* (New York: Twelve).
- Jansson, Lina and Jonathan Tallant forthcoming, 'Quantitative Parsimony: Probably for the Better', *British Journal for the Philosophy of Science*.
- Lewis, David 1973, *Counterfactuals* (Oxford: Basil-Blackwell).
- Morgan, C. Lloyd 1894, *An Introduction to Comparative Psychology* (London: Walter Scott).
- Nolan, Daniel 1997, 'Quantitative Parsimony', *The British Journal for the Philosophy of Science* 48, pp. 329–343.
- Povinelli, Daniel J. and Jennifer Vonk 2004, 'We Don't Need a Microscope to Explore the Chimpanzee's Mind', *Mind and Language* 19 pp. 1–28.
- Sober, Elliott 1975, *Simplicity* (Oxford University Press).
- 1981, 'The Principle of Parsimony', *The British Journal for the Philosophy of Science* 32, pp. 145–156.
- 1988, *Reconstructing the Past – Parsimony, Evidence, and Inference* (Cambridge, MA: MIT Press).
- 1990, 'Let's Razor Ockham's Razor', in D. Knowles (ed.), *Explanation and Its Limits* (Cambridge University Press), pp. 73–94.
- 2001, 'What is the Problem of Simplicity?', in A. Zellner, H. Keuzenkamp, and M. McAleer (eds.), *Simplicity, Inference and Modelling: Keeping It Sophisticatedly Simple* (Cambridge: Cambridge University Press), pp. 13–31.
- 2008, *Evidence and Evolution: The Logic behind the Science* (Cambridge University Press).
- 2009, 'Parsimony Arguments in Science and Philosophy — A Test Case for Naturalism', *Proceedings and Addresses of the American Philosophical Association* 83(2), pp. 117–155.
- Swinburne, Richard 2004, *The Existence of God* 2nd Edition (Clarendon Press).
- 2009, *Simplicity as Evidence of Truth* (Milwaukee: Marquette University Press).

Tomasello, Michael and Josep Call 2006, 'Do Chimpanzees Know what Others See, or Only What They are Looking at?', in S. Hurley and M. Nudds (eds.), *Rational Animals?* (New York, NY: Oxford University Press), pp. 371–384.

Whiten, Andrew 1996, 'When does Smart Behavior-Reading become Mind-Reading?', in P. Carruthers and P. Smith (eds.), *Theories of Theories of Mind* (Cambridge University Press), pp. 277–292.

Texas Tech University
joel.velasco@ttu.edu

JOEL VELASCO

doi:10.1093/mind/fzx031

Advance Access publication 16 October 2017

Inner Virtue, by Nicolas Bommarito. New York: Oxford University Press, 2018. Pp. 208.

Suppose I told you that the person you consider your best friend often dwells on your faults in his own mind; while he behaves in a warm and affectionate manner when the two of you are together, privately, he ruminates on his advantages over you. He likes to compare himself to you because he finds the comparisons flattering. He does not see you as his equal and the thought that he is better pleases him.

This is a chilling thought. But why is it chilling? Part of the reason has to do with the connection between inner states and actions. Perhaps you think that your friend will, sooner or later, do something contrary to the spirit of friendship. Secret feelings and desires may surface and shape behaviour. But this is hardly the whole story. For we could easily imagine that the privately held negative attitudes will *not* affect action. Maybe your friend is deceased and you only learn from his diary how he saw you. If that happens, you will probably feel hurt, despite the fact that the dead can do nothing to you. Moreover, you are likely to revise your judgment of your friend's character – you will come to think that he was not as good as you thought.

Would such revisions of character assessment based solely on inner states be warranted? In *Inner Virtue*, Bommarito argues that the answer is 'yes'. He discusses three types of inner attitudes: pleasure, emotion, and attention. What we delight in, are moved by, and attend to can make us morally better or morally worse. This is because, in Bommarito's view, pleasure, emotion, and attention have a connection to moral concern, understood as concern for moral goods.

On Bommarito's reckoning, moral philosophers have tended to focus almost exclusively on action and agency. When it comes to inner states, only those that bear on action – such as motives and intentions – have received discussion. But there is more to being a person, and to being a good person, than our actions and their motives: we have inner lives that