The Geometry of Desert, by Shelly Kagan. Oxford: Oxford University Press, 2012. Pp. xvii + 656. H/b L47.99, p/b L25.99.

Most philosophy books, it's fair to say, contain more footnotes than graphs. By this standard Shelly Kagan's *The Geometry of Desert* is very much an outlier. Its subject is moral desert, or the intrinsic value of people's getting what they morally deserve. More specifically, it examines the structure of this value, or what determines exactly how good it is if people with specified degrees of moral virtue enjoy specified amounts of happiness. Kagan illustrates the relevant possibilities with graphs, and one of his aims is to demonstrate the fruitfulness of graphs as a tool in ethical theory. But another is to show that many more structures are plausible for desert than we might initially think, so this value is more complex than we might think. The result is a book with 203 graphs. They perspicuously represent the views he considers and also help him discover them. To complete a graph you often have to decide exactly where a line will go or what shape it will have. This choice can reflect a moral issue you wouldn't otherwise have thought of but that, once identified, is interesting and important. Part of the fruitfulness of Kagan's graphs is that they point to many such issues.

At the same time, his discussion is almost entirely self-contained, with few references to other philosophical treatments of desert, either classical or present-day. The result is that a book of 641 pages has only fifteen footnotes citing other published work. Some may wish Kagan did more to relate his ideas to others' and explicitly address their arguments. But a wise reader will welcome his firm focus on his own line of argument and the unified treatment that results, free of scholarly digressions. *The Geometry of Desert* is an exciting ride through largely uncharted philosophical territory that repeatedly identifies new issues and proposes new concepts to deal

with them. The book is an analytical tour de force and a splendid exemplar of constructive philosophy; it's also extremely clearly written and a pleasure to read.

Like others, Kagan distinguishes two forms of desert: noncomparative and comparative. Noncomparative desert asks whether an individual considered on his own is getting what he deserves, regardless of what anyone else has or deserves. Comparative desert, by contrast, looks at a pattern of distribution across individuals and asks whether it fits the pattern of their deserts. The first half of the book discusses noncomparative desert and the second half comparative.

Noncomparative desert says it's good if a virtuous person is happy and, if it has a retributive side, that it's good if a vicious person suffers. Kagan thinks it says, more specifically, that for any specified degree of virtue there's a specific amount of happiness a person with that degree of virtue ideally deserves, or that it's best in desert terms for him to have. If he has a little more or a little less than that amount, that's good though not ideally so; if he has a lot more or less, that can be bad.

He represents this view on graphs whose x-axes measure a person's happiness and whose y-axes measure the value "from the point of view of desert" of her enjoying that happiness. That a person is happy is good, and that she's virtuous may also be good. Noncomparative desert finds a further value in the combination of the two in one life, or in the fact that it's a virtuous person who's happy. On Kagan's graphs it generates a series of "desert mountains," as in Figure 1. Here each "peak" represents the greatest desert-value possible for a given person, and involves her having exactly the happiness she ideally deserves; greater or lesser amounts are on the sides of the mountain. And since more virtuous people ideally deserve more happiness, the sequence of mountains from left to right represents the noncomparative deserts of progressively more virtuous individuals. Kagan then considers a series of detailed questions about the shapes and

mutual relations of these mountains and the moral issues they reflect.

His account of comparative desert piggybacks on his account of noncomparative desert. He assumes, plausibly, that if noncomparative desert is completely satisfied, with each person at her peak, comparative desert is also completely satisfied. But imagine that one person A has either somewhat less or somewhat more happiness than she ideally deserves. She's at a point some distance down the y-axis from her peak, where this distance measures the loss in noncomparative value of her having what she does. According to Kagan's "Y gap" view, what comparative desert requires for a second person B is that he fall short in the same way, by likewise having too little or too much, and by the same amount of desert-value. If A's having the happiness she does means a loss of 5 units of noncomparative value compared to her peak, B should have whatever happiness makes for an equal loss of 5 units from his peak, so what comparative desert seeks to equalize is "offence against noncomparative desert" measured in y-axis terms. Kagan again discusses more detailed questions about this view, such as how it measures degrees of comparative value in distributions whose Y gaps aren't equal.

The remainder of this review will try to situate his treatment in relation to others by comparing it to a more traditional view that associates desert with proportionality. Aristotle called distributive justice "a species of the proportionate" and said it's satisfied when the ratio of what A gets to what B gets equals the ratio of A's merit to B's merit; Kant and Ross, too, wanted happiness proportioned to virtue. Their view can be incorporated in noncomparative desert if the latter satisfies a proportionality condition whereby, for any fixed amount of happiness, the distribution of it that's best, or results in the most noncomparative value, is proportional. And comparative desert can value proportional distributions directly. How do Kagan's accounts relate to these two?

Though he doesn't mention proportionality in connection with noncomparative desert, a view embodying it can be represented on his graphs. It too generates desert mountains with peaks, but the proportionality condition forces certain choices about them that he leaves open. (He's often more interested in displaying options than in deciding between them.) As against the "straight desert" represented in Figure 1 it requires what he calls "curved desert," where the sides of the mountains are rounded, their slopes flattening as they approach the peak. This has the attractive implication that, if A and B are equally virtuous but B has much less happiness than he ideally deserves while A has only a little less, it's better in desert terms to give a unit of happiness to B than to A. (Kagan is attracted to curved desert but sets it aside because it would complicate his discussion of other issues.) A related condition requires the peaks for more virtuous people to be higher, or further up the y-axis, so if A is twice as virtuous as B, the noncomparative value of her getting the happiness she deserves is twice that of B's getting his ideal happiness. This too seems attractive. If you could give either Nelson Mandela or a slightly virtuous person the happiness he ideally deserves, wouldn't it be better, even just in desert terms, to give it to Mandela? Likewise if you could give the deserved punishment to a great or to a minor sinner. Kagan leaves this possibility, too, open but, surprisingly, thinks most people will prefer a view with peaks of equal height. Finally, the proportionality condition requires the desert mountains to be isomorphic, so they all have the same shape and orientation, though they may be proportionally larger, or higher and wider, for more virtuous people. But Kagan rejects isomorphism in his discussion of what he calls "bell motion."

Consider whether the slope of a mountain should be steeper on its right side or on its left, which reflects the normative question whether it's worse, or a greater loss of desert-value, to have more happiness than you ideally deserve or to have less. (This is another novel issue the

graphs suggest.) Kagan thinks there's no uniform answer. His preferred view is that if you're very virtuous, it's worse to get less than you deserve and less bad to get more, whereas if you're only a little virtuous it's worse to get more. The orientations of his desert mountains therefore change. Toward the right of the graph, where they represent desert for very virtuous people, their base points out toward the right. Further left, and especially if they represent desert for vicious people, their base points toward the left. The mountains swing in a clockwise direction as they move from right to left, mimicking the motion of a bell (Figure 2).

A proportionality view of noncomparative desert agrees that it's less bad for more virtuous people to be a unit of happiness beyond their peak, since the slope on a larger mountain changes more slowly. But for the same reason it denies that it's worse for the more virtuous to be a unit short of their peak; instead it's less bad. I'm not sure which view here, the proportionality one or Kagan's, is more intuitive. But bell motion as represented in Figure 2 also implies that if a vicious person suffers more than he ideally deserves, that's less bad than if he suffers less than he deserves; and if he's very vicious, his suffering more than he deserves is hardly bad at all. I think many will say, to the contrary, that punishing even a bad person too much is a greater injustice than punishing him too little. Kagan may reply that we're influenced here by the evil of the suffering as suffering rather than as deserved, but I think many will maintain it just as a view about justice. It's one a proportionality view can endorse but his preferred version of bell motion rejects.

Kagan's view of noncomparative desert doesn't differ that fundamentally from one governed by proportionality. His graphs allow illuminating representations of that view, he raises new issues its defenders must address, and he differs from it only on some points of detail. His treatment of comparative desert, by contrast, is a much more radical departure.

He starts by explicitly rejecting a proportionality or "ratio" view of comparative desert, on the basis of some original and powerful objections. Imagine that A is twice as virtuous as B. A proportionality view says it's best if A is twice as happy as B, but what if B is unalterably suffering 10 units of pain? Does the view now demand that A suffer 20 units of pain, so she again has twice as much? That's absurd. Or imagine that A's degree of virtue is zero while B's is positive. If A has zero happiness, as we can imagine she ideally deserves, it seems proportionality is satisfied no matter what B has, e.g. 10 units of happiness or a million. Kagan thinks that, too, is unacceptable. In a final case A is vicious, B is virtuous, and A suffers less pain than he ideally deserves. Does proportionality demand that B now enjoy less happiness than he deserves? He again finds that hard to believe.

A proportionality view of comparative desert can avoid some of these objections if it's formulated in ways its classical defenders didn't envisage. It can say that there's not only an intrinsically best division between twice-as-virtuous A and half-as-virtuous B but also a worst one: it's the contrary-to-proportional division in which A suffers twice as much pain as B. It then doesn't applaud this division, as Kagan assumes the classical view does, but condemns it. (The values of other divisions depend on where they fall between the best and the worst.) The view can also deny that there's a single best division between vicious A and virtuous B. It can value the proportioning of happiness among the virtuous and of pain among the vicious but say nothing specific about mixed cases. Nor is this implausible. Do we think we should proportion Nelson Mandela's happiness to Idi Amin's suffering, or the circles of heaven to those of hell? These responses reject what Kagan calls the "optimistic" assumption that given fixed degrees of virtue and vice for two people and a fixed amount of happiness or suffering for one of them, it's always possible to assign happiness or suffering to the other in a way that completely satisfies

comparative desert. But it's not clear that optimism is a compelling requirement on comparative desert or that anything in the classical writers commits them to it.

Kagan's Y gap view of comparative desert takes a completely different line. By building on noncomparative desert it can accept optimism, since given a certain Y gap for A it's always possible to give the same gap to B; he finds this attractive. It also has plausible implications in the cases he raises against the proportionality view. Is it then the clearly best comparative view?

The Y gap view is original, even brilliantly so, and certainly a serious contender. But it has some questionable implications of its own. If the peaks for more virtuous people are higher up the y-axis, as they can be even apart from proportionality, its ideal distributions can greatly favour the more virtuous. Imagine that B is five times as virtuous as A and has a peak five times as high, and that A unalterably has half the happiness she ideally deserves. Giving B the same absolute Y gap, as Kagan's view recommends, means giving him 90% of what he ideally deserves, or nine times as much happiness as A. (I here assume straight desert without bell motion; adding the latter favours B even more.) Those with proportionality intuitions will balk at this claim. The Y gap view also says that if two people who are equally virtuous suffer equally intense pain, there's something good in the situation. A proportionality view that rejects optimism can deny that.

Kagan cares that the Y gap view have intuitive appeal in itself, apart from its implications, and the idea of equalizing offence against noncomparative desert gives it some. (Whether this is as much as a proportionality view has on its own is less clear.) But the view can seem at once too close to and too far from noncomparative desert.

Because it's formulated in terms of noncomparative desert, the Y gap view makes it impossible to value only comparative and not noncomparative desert, as Aristotle seems to have

done. Is that reasonable? It also makes comparative desert hostage to the details of noncomparative desert, its verdicts changing with changes in, e.g., the degree of bell motion. Is it not more independent? At the same time Kagan's composite view lacks an attractive feature found in ones governed by proportionality. Their noncomparative and comparative parts always agree about the best division of a fixed amount of happiness – it's proportional – and differ about that only when the amount can be increased or decreased. (Compare how, within egalitarianism, equality and priority views agree whenever the total to be distributed is fixed.) Nothing in Kagan's view guarantees this overlap. Even given a fixed total happiness, his noncomparative and comparative principles can disagree about its optimal division. Though none of these objections is decisive against the Y gap view, they show that, like the alternatives, it has demerits as well as merits.

The Geometry of Desert is essential reading for those who value moral desert, but it also has a great deal to offer those who don't. Kagan doesn't argue directly for the value of desert, focussing instead on how we should understand it if we do value it. But if we find the structural questions he raises about it real ones, with some intuitive appeal on both sides, that's indirect evidence that desert really is of value. If it weren't, how could choices about it engage us? His discussions also make important methodological points. One of his aims is to show that desert is more complex than we might think, and in that he succeeds spectacularly; there's more to desert than has been dreamt of in previous philosophy. But the phenomenon is more widespread. Our initial thoughts about many moral considerations – equality, self-defence, thresholds, virtue – take them to be simpler than on closer examination they turn out to be. Kagan shows dramatically, in his one instance, how much careful moral analysis can uncover. Finally, there's the great value he shows in graphs, which are relevant to many moral topics other than desert.

Though all his issues could in principle be raised just verbally, they're vastly easier to see and discuss using graphs. He shows, we can say, that a picture is worth a thousand footnotes.

Thomas Hurka

Department of Philosophy

University of Toronto

170 St. George St.

Toronto, ON M5R 2M8

CANADA

tom.hurka@utoronto.ca