

Comparative Desert

SHELLY KAGAN

1. An adequate theory of (moral) desert will need to include not only noncomparative principles, but comparative ones as well. Accordingly, my central concern in this paper will be to evaluate some candidates for the basic comparative principle.

I am going to begin, however, by sketching a few of the main elements of the theory of noncomparative desert, and then arguing, briefly, that this noncomparative theory needs to be supplemented with a comparative theory. As we shall see, although it is fairly easy to evoke intuitions that point toward the need for a comparative principle of some sort, it is not at all obvious what that comparative principle should look like.

Let me start, then, with some rather dogmatic pronouncements about the nature of desert.¹ In other contexts I would be prepared to defend these claims at greater length, and there are a great number of details and complications that I will be passing over with little or no comment; but my concern here is simply to make clear the outlines of the machinery upon which I will be drawing later.

I take it that people vary in terms of how deserving they are; and while it is controversial what precisely the basis of being more or less deserving is, it is natural to talk in terms of varying levels of virtue, where the more virtuous are more deserving and the less virtuous are less deserving. I take it as well that it is a good thing—other things being equal—if people get what they deserve. No doubt there are a variety of desirable effects that typically ensue when people get what they deserve, but I believe that there is intrinsic value in people getting what they deserve as well, and not merely instrumental value, and my concern here is solely with the intrinsic value of people getting what they deserve.

¹ Here, and in all that follows, I draw from a considerably larger work in progress, *The Geometry of Desert*.

But what is it that people deserve? Presumably, when people vary in terms of how deserving they are, they vary in terms of what they deserve: the more deserving individuals deserve *more* of something. And while this is controversial as well, I am simply going to assume that what they deserve more of is well-being. The more virtuous deserve to be *better off* than the less virtuous, who in turn deserve to be better off than the vicious. Some would go further, of course, and insist that if someone is sufficiently vicious they may actually deserve to suffer. I have a great deal of sympathy for this retributivist view, but it shouldn't be particularly important in most of what follows.

I am however going to assume that for each person there is a *particular* level of well-being that is deserved by that person—what they “absolutely” deserve (though it will, of course, vary from individual to individual). After all, even those retributivists who hold that some people deserve to suffer believe that for any given individual there is a limit to how badly off that person should be. And although it is more controversial, I believe that something similar is true even for those who are virtuous—even to remarkable degrees. The virtuous doubtless deserve to be well off, maybe extremely well off. But I believe (perhaps somewhat cynically) that no matter how virtuous a person is, there is still some level of well-being that is the particular level of well-being absolutely deserved by that person. For any individual at all, then, it is possible for that person to have less than they deserve (that is, to have a lower level of well-being than they deserve), but it is also possible for that person to have more than they deserve (that is, a higher level of well-being than they deserve).

It is natural to assume that if someone has precisely what they absolutely deserve, this is a good state of affairs, other things being equal. But of course other things may not be equal, since any number of factors may affect the overall intrinsic value of a given state of affairs. Since our concern here is only to map the contribution to intrinsic value made by the fit or lack of fit between what people have and what they deserve (that is, the extent to which people are getting the particular level of well-being they deserve), let us hereafter restrict our attention to what I will call “goodness from the standpoint of desert”—that is, the direct contribution to the intrinsic value of a state of affairs due to desert (rather than equality, say, or the intrinsic value of virtue). This will allow us to drop the constant qualification “other things being equal.” We can then say, more simply, that it is natural to assume that if someone has what they absolutely deserve, this is a good state of affairs.

Suppose, however, that they have more or less than they deserve. This will be a less good state of affairs (from the standpoint of desert). There will be a drop off in value, since by hypothesis things would be even better if the person had exactly what they deserved, rather than too much or too little. And the further the person is from the level they absolutely deserve, the greater the drop off. Indeed, if the person is sufficiently far from having what

they deserve—whether they have far too little or far too much—this may be sufficient to create an intrinsically bad state of affairs.

I find it helpful to think about these issues in graphic terms. Suppose we let the X axis represent a given individual's (possible or actual) level of well-being, where points to the right of the origin are lives worth living, and points to the left of the origin are lives worth not living. Then the Y axis can represent goodness from the standpoint of desert, where points north of the origin are intrinsically good with regard to desert, and points south of the origin are intrinsically bad. (Since I am here sketching the theory of noncomparative desert, it would of course be more precise to say that the Y axis represents goodness from the standpoint of *noncomparative* desert; but I'll keep the qualification implicit until it's relevant.)

Figure 4.1 shows the individual desert graph for a possible individual. The “peak” of the mountain represents the level of well-being that the individual absolutely deserves. If this is what the person has, this is optimal. But if he has too much, or too little, there is a drop in intrinsic value, represented by the eastern and western slopes, respectively. Other individuals, who differ from this first, by being more or less virtuous, would have similar looking desert graphs, but the mountains would be shifted to the east or to the west.

Not surprisingly, there are many complications that a fuller discussion would need to address. Let me mention just three of them briefly. First, while I do believe that everyone, no matter how virtuous or vicious, will have a desert graph with the same basic mountain-like shape, I do not in fact think that the differences in these graphs are limited to the fact that the more virtuous an individual the further the mountain's shift to the east. On the contrary, I believe that the slopes on the eastern and western sides of the mountains will vary as well. In particular, I believe that the slope of the western side of the mountain grows steeper with more virtuous individuals (so that if one must




Fig. 4.1. A representative individual desert graph.

leave either a virtuous individual or a vicious individual a certain amount short of their respective peaks, it is worse to shortchange the virtuous individual). And I believe that the slope of the eastern side of the mountain grows gentler with more virtuous individuals (so that if one must overcompensate either a virtuous individual or a vicious individual by a certain amount, it is less bad to overcompensate the virtuous individual). Putting these two points together, we find that the mountain actually swings like a bell (if we think of the peak as fixed), to the left for vicious individuals and to the right for virtuous ones. Not everyone accepts *bell motion*, but I find it plausible.

Second, while I have drawn the western and eastern slopes as straight lines, it is plausible to think that in fact they may be curved, so that the further one is from one's peak (whether to the east or to the west) the steeper the drop off, and the greater the significance of each further alteration in well-being. But incorporating this feature—*curved desert*—would complicate the discussion, and the graphs, without affecting any of the main points, so I will disregard it in what follows.

The final potential complication is this. I have already noted that the location of the peaks will vary along the X axis (so that more virtuous individuals will have peaks further to the east). But what should we say about variation along the Y axis? One possibility of course is that all peaks have the *same* (positive) Y coordinate, so that the *skyline* (the line consisting solely of the various possible peaks) is a straight line parallel to the X axis. This would represent the claim that while people vary in terms of what they deserve, it is nonetheless equally *important* that everyone get what they deserve. But there are other possibilities as well. For example, it might be that the skyline should have a V shape, with the bottom of the V at the origin. This would represent the rival claim that when people are very deserving—whether of good or of ill—it is *more* important that they get what they deserve. This is an important debate, but I won't take a stand on it here.

While these various complications are important for developing the theory of noncomparative desert, luckily in most of the discussion that follows they won't concern us. I mention them here only because there will be a few points toward the end where they are indeed relevant (two of them, at least) for comparing some alternative accounts of comparative desert.

2. Now the theory that I have been sketching so far is essentially *noncomparative*. For any given individual, to determine the intrinsic value of his being at a given level of well-being I need only consider his individual desert line. By determining whether he has more or less than he deserves (whether he is to the east or to the west of his peak) I learn whether I can make the situation better or worse from the point of view of desert by changing his level of well-being. But I do not need to consider what other people deserve, or how

well off they may be. That's what I mean in calling the theory noncomparative. Of course, we can certainly compound and compare these various noncomparative judgments. In this way I might learn, for example, that I can do more good by aiding one individual a certain amount rather than another. But the information being compounded—the facts about what each person deserves and how much good it would do to alter his level of well-being—this information is, for all that, essentially noncomparative in nature.

A complete theory of desert, however, will also have to include essentially comparative elements as well. For although it matters whether I get what I (absolutely) deserve, this is not all that matters. It also matters how I am doing *compared* to you, in light of how (noncomparatively) deserving we are. That is the basic idea of *comparative* desert. Thus, for example, if I am just as deserving as you are (just as virtuous), then I should be doing as well as you (no matter how well you are doing). Similarly, if you are more virtuous than I am, then you should be better off than I am (no matter how well off I am). Of course these things may not be true all things considered. Indeed, even if we restrict our attention to desert, these comparative considerations will often be opposed by noncomparative considerations (since noncomparative desert will oppose giving either of us more than what we absolutely deserve). But the point for now is simply that most of us do feel the pull of these comparative considerations as well. When I am as virtuous as you, then I should be doing as well as you, no matter how well you are doing. If I am not, then there is something to be said in favor of improving my lot to bring me up to where you are—regardless of where you are. That is the claim of comparative desert. It is an essentially *comparative* claim, since it is concerned essentially with comparing our levels of well-being—in light of how (noncomparatively) deserving we are.

Suppose, for example, that A and B are equally deserving—have the same peak—and A is at the peak but B is beyond it (at a higher level of well-being). Imagine that we cannot alter B's level of well-being, but we could improve A's. Should we?

Noncomparative desert says no, since raising A's level of well-being only serves to give him more than he deserves, and this—from the standpoint of noncomparative desert—can only make things worse. The fact that B is beyond her peak is bad, but moving A beyond his peak as well doesn't make things better.

From the standpoint of comparative desert, however, there is indeed something to be said in favor of moving A. After all, A is just as deserving as B, and so deserves to be as well off as A is. Putting A at the same level as B is thus an improvement from the perspective of comparative desert.

Or imagine that A is *more* deserving than B, but B has more than A. Even if A is already at his peak, isn't there something to be said in favor of improving his lot even more, so that he has more than B? Noncomparative desert says

no: B is beyond her peak and that is bad, but moving A beyond his peak will not improve the situation at all. But comparative desert says yes: A is more deserving than B, and so there is something to be said in favor of increasing A's level of well-being, so that he ends up better off than B.

In cases like this we have a conflict between comparative and noncomparative desert. Accordingly, those of us who want to incorporate both aspects into a complete theory of desert will eventually need to work out a tradeoff schedule, so that we know which has more weight in such cases of conflict. But that will not be my concern here, for the question is premature until we have developed an adequate understanding of the two competing principles. My immediate concern has simply been to argue for the intuitive attractiveness of such comparative claims. Since they cannot be captured from the standpoint of noncomparative desert, we must supplement such a theory with a comparative component as well.

Of course, some people may prefer to couch these comparative claims in terms of "fairness" rather than "desert." They may happily concede that the comparative value I have just been pointing to is a genuine one, well worth exploring, but insist that desert proper is limited to noncomparative desert alone; what I have just been describing as comparative desert is, rather, a matter of fairness. (If I am as deserving as you, then it isn't fair if you have more than me; if you are more deserving than me, it isn't fair if you aren't better off than I am. And so forth.) But as far as I can see nothing important turns on this dispute. Judgments about fairness, after all, can be sensitive to many kinds of differences (or similarities), and all that is important for my purposes is that we recognize that one set of fairness claims turns upon judgments about how I am doing compared to you, in light of how deserving we both are. One could, I suppose, call this "desert-sensitive fairness." I prefer to call it "comparative desert." The important point, for our purposes, is to recognize that we do indeed want some account of comparative desert (whatever we call it).

Unfortunately, while I think that most of us can fairly readily see the value of comparative desert—that there is at least some kind of intrinsic value in satisfying the claims of comparative desert—it is far less obvious what precisely the requisite comparative principle would look like.

The basic idea, of course, is clear: comparative desert demands that my level of well-being bear a certain relation to your level of well-being, where this precise relation is itself a function of how our levels of virtue compare. But what, exactly, is the relevant relation (or the relevant function)? This is a matter of some dispute.

For example, suppose that your peak is 20 (units of well-being) and mine is 10. Imagine, however, that I am actually at 30 and there is nothing that can be done about this. But we do have the ability to alter *your* level of well-being. Where then should you be to satisfy comparative desert?

We have, of course, already noted the intuition that you should have *more* than me (since you are more virtuous than I am). But how much more? If you are at 40 is that enough? Should you be at 60? More? Less? Unfortunately, until we have the correct comparative principle, we are not yet in a position to settle most questions like this. (Of course, if you and I are equally deserving then it does seem clear that comparative desert is perfectly satisfied when we are at the same level of well-being—whatever that may be. But this is a special case.)

3. Presumably, there will be various plausible constraints on an adequate theory of comparative desert. I have already noted two (if we are equally deserving, we should be equally well off; if you are more deserving, you should be better off), but let me quickly mention one more. I find the following claim extremely attractive: when noncomparative desert is perfectly satisfied, comparative desert is perfectly satisfied as well. That is to say, if everyone is exactly at their peak (so that the situation is optimal from the standpoint of noncomparative desert), then the correct comparative relation—whatever it is—obtains as well (so that the situation is optimal from the standpoint of comparative desert).

This is not to say, of course, that the *only* way to satisfy comparative desert is to have everyone at their peaks. Since comparative desert is concerned with relations, rather than absolute levels, presumably it will often be possible to satisfy it even though the demands of noncomparative desert remain unsatisfied. But although having people at their peaks is not necessary to satisfy comparative desert, it is, I believe, sufficient. Whatever it is that the demands of comparative desert come to, they are perfectly satisfied when noncomparative desert is perfectly satisfied as well.

Admittedly, I have no argument for this claim, and so offer it only as a conjecture.² If it is not true, it seems to me, then the two parts of the theory of desert—comparative and noncomparative—float free of each other in an implausible and philosophically unsatisfying way. So I am simply going to assume that it is true.

Obviously enough, accepting our conjecture does narrow the field of potential comparative principles (since it rules out all proposals according to which comparative desert might be unsatisfied even though everyone is at their peak). Unfortunately, it doesn't suffice to allow us to settle upon the correct one. By hypothesis, of course, if you and I are both at our peaks the relevant comparative relation obtains. But knowing this fact isn't enough to determine what that relation is, and so we are not yet in a position to say how to satisfy comparative desert in those cases where one of us is *not* at our peak.

² Because of the attractive consonance, I like to think of this conjecture as "Kagan's conjecture."

Consider again the case where my own peak is 10 but I am in fact at 30, and your peak is 20. Presumably there is *some* level of well-being (greater than 30) such that if you are in fact at that level comparative desert will be perfectly satisfied—despite the fact that neither of us will be at our peaks. But what is that level? Of course, given our conjecture, we know that if I *were* indeed at 10 and you at 20, comparative desert would be satisfied. But what of it? I am *not* at 10. We are not yet in a position to tell which of the relations that obtain in that ideal case is the relevant relation; so we don't know what to look for in other situations.

For example, some might hold that what is important about this case is that your peak is 10 units higher than my own. Accordingly, they might claim, comparative desert will be satisfied whenever you are 10 units better off than I am. Thus, since my own well-being is fixed at 30, you should be at 40. Others, however, would insist that what matters is not the absolute difference between our peaks, but rather their ratio. Since your peak is 20 while mine is 10, you are twice as deserving as I am. Accordingly, it might be claimed, comparative desert will be satisfied whenever you are twice as well off as I am. Thus, since my own well-being is fixed at 30, you should be at 60 (rather than 40). And, obviously enough, various other proposals could be offered as well.

So our situation is not yet significantly altered. It is plausible to think that we need a comparative principle, if we are to have an adequate theory of desert. But it is far from clear what the correct comparative principle is.

4. I imagine that the most widely accepted candidate for the comparative principle is the *ratio view*. According to this view, comparative desert is satisfied when my level of well-being stands to your level of well-being as my level of virtue stands to your level of virtue. Given the natural assumption that peaks are themselves fixed in proportion to virtue, this comes to the same thing as the claim that the ratio between our levels of well-being should be the same as the ratio between our peaks.³ Thus, if you are twice as virtuous as me (if your peak is twice as far to the right of the origin as mine) you should be twice as well off. If I am one third as virtuous as you (if my peak is one third as great) then I should be one third as well off.

The ratio view has an undeniably attractive ring to it, and what is more it has an impeccable pedigree, since it seems to have Aristotle's backing as well.⁴ But it is incorrect.

³ Of course this natural assumption—that peaks are fixed in proportion to virtue—could be challenged, in which case we will need to distinguish between ratio views that fix the relevant ratio in terms of virtue and those that fix it in terms of peaks. But for our purposes we need not distinguish between these views, since both are subject to the sorts of objections I will be raising here. For simplicity, therefore, let us suppose that the natural assumption is correct.

⁴ See the *Nicomachean Ethics*, 1131^a10–1131^b24, where Aristotle describes comparative justice in terms of the ratio view. Admittedly, Aristotle is there making a claim about justice, rather than

Various cases provide difficulties of different sorts for the ratio view. Let me quickly mention three. Suppose that A's peak is at 10, and B's is at 20. B is twice as virtuous as A, and should, therefore, be twice as well off. It is, of course, a straightforward matter to apply the ratio view to those cases where A's actual level of well-being is a positive real number. If A is at 3, B should be at 6; if A is at 1200, B should be at 2400. But it is not so clear what advocates of the ratio view should say if A's actual level of well-being is a *negative* number (a life not worth living). Suppose, for example, that although A's peak is at 10, we can do nothing about the fact that his level of well-being is at -10. Where should B be placed so as to satisfy comparative desert?

Since B is twice as virtuous as A, advocates of the ratio view seem committed to the view that B's level of well-being should be twice that of A's. But this seems to mean that if A's level is fixed at -10, the ratio view demands that B be placed at -20! (For $2 \times -10 = -20$. More precisely, since A's peak is 10 and B's is 20, B's level of well-being should stand to A's as 20 stands to 10. But $-20/-10 = 20/10$. So if A is fixed at -10, B should be at -20.)

But this, I take it, is absurd. B is by hypothesis a more virtuous individual than A. It seems clear that in a case like this comparative desert demands that B be better off than A. (This was, in fact, one of the original intuitions that I used to motivate the need for a comparative principle in the first place.) It is quite unacceptable to claim that comparative desert insists that the more virtuous B should be placed at a level significantly *worse* than A! But this does seem to be the most straightforward way of understanding the implications of the ratio view. So understood, then, the ratio view is simply unacceptable.

There is, however, an alternative interpretation of the ratio view that escapes this objection. Perhaps when we switch over to the negative case, we should *reverse* the ratio. B is twice as virtuous and so should be twice as well off. When we are dealing with positive levels of well-being, this means twice as great a level (that is, two times A's level of well-being). But when we are dealing with negative levels of well-being, it might be suggested, this means *half* as much. Since B is twice as virtuous as A, if A must suffer B should suffer only half as much. Thus, if A is fixed at -10, comparative desert is satisfied when B is located, not at -20, but at -5.

This alternative interpretation avoids the absurd conclusion that B should be placed at a lower level than A, despite being more virtuous. In this case, at least, it accommodates the intuition that comparative desert demands that the more virtuous be better off than the less virtuous. And it does this in a way that still preserves the central significance of the ratio of levels of virtue, albeit by reversing the ratio in certain cases. It thus seems superior to the original

desert *per se*, but given the close connection typically thought to exist between the two, it is natural to ascribe to Aristotle the belief that the correct principle of comparative desert will be stated in terms of the ratio view as well.

interpretation of the ratio view. (Similarly, if we have a case where both A and B have negative peaks, but one is fixed at a positive level of well-being, we can avoid absurdity if we again reverse the ratio, guaranteeing that the more vicious individual deserves a lower level of well-being.)

But both versions of the ratio view remain subject to further difficulties. Several troubling cases involve having one of the relevant values be set to zero. It is, I suppose, a familiar point that ratio views of various kinds typically have difficulty handling cases involving zero, but the familiarity of the point doesn't make it any less telling. For the sake of brevity, let me here mention only one case of this sort. Suppose that A's peak is 0 (that is, what he absolutely deserves is a life neither worth living nor worth not living), and B's peak is 10. And imagine that A is in point of fact at 0. Where should B be placed?

Since A is already at his peak, given our earlier conjecture we know that if B is placed at her peak as well, then comparative desert will be satisfied. And sure enough, both versions of the ratio view accommodate this point, since if B is placed at 10, then the ratio of levels of well-being will clearly be the same as the ratio of peaks, that is (let us say) 0 to 10 (A to B). Unfortunately, it seems arguable that the ratio view will also be satisfied *regardless* of where B is placed, so long as A remains at 0. For example, even if B is placed at 100, it might be argued that the relevant ratio has been maintained. (A should have 0 units of well-being for every 10 units that B has. And he does, even when B is at 100, provided that A is at 0.) But this seems absurd as well. That is, it seems quite unacceptable to suggest that comparative desert is *indifferent* in this case as to how we place B. (On the contrary, it seems plausible to hold that if A is at his peak, as in this case, the *only* location for B that satisfies comparative desert is to have B at her peak as well. But it is not at all clear that the ratio view can satisfy this constraint.)

Indeed, it seems that the ratio view will be satisfied even if B is given a *negative* level of well-being (since A will still have 0 units of well-being for every 10 that B has). And this, of course, would once again violate our initial intuition that comparative desert insists that the more virtuous (B) should be better off than the less virtuous (A).

Presumably there are ways for advocates of the ratio view to try to avoid this result. They might, for example, insist that when A's peak is 0, the relevant ratio is undefined (since division by zero is undefined, and so 10/0 is undefined). But this has its own costs, since it now leaves the ratio view unable to acknowledge the plausible conjecture that if both A and B are placed at their peaks then comparative desert is satisfied. (For if the relevant ratio is undefined, how can it make sense to say that when both A and B are at their peaks the levels of well-being stand in the relevant ratio?)

Let me turn now to a third type of case that is problematic for ratio views. Imagine that A's peak is -10 (he deserves to suffer somewhat) and B's peak is +20 (she deserves to be well off). Here it seems clear that the relevant ratio

must be -10 to +20 (A to B). That is, for every -10 units of well-being had by A, B must have +20 units of well-being. Equivalently, for each *negative* unit of well-being had by A, B must have *two positive* units of well-being. But now imagine that A's actual level of well-being is fixed at -5. Where should B be placed to satisfy comparative desert?⁵

It seems that advocates of the ratio view must claim that B should be placed at +10. This level, and only this level, maintains the appropriate ratios (since $+10/-5 = +20/-10$); this level and only this level gives B two positive units of well-being for every one of A's negative units.

But this result is completely unacceptable. Think about what is being said here. A is a vicious individual and noncomparatively deserves to be at -10. As it happens, he is at -5, a higher level of well-being than he absolutely deserves. But there is nothing we can do about that: A unavoidably has more than his peak. What then does comparative desert tell us to do in this case? According to the ratio view, it tells us to take the more virtuous individual, B, and leave her *lower* than her peak!

This is absurd. Certainly comparative desert cannot instruct us to move a more virtuous person *below* her peak in response to a less virtuous person being above his peak. If the less virtuous are getting more than they absolutely deserve, surely comparative desert demands that the more virtuous should be getting more than they absolutely deserve as well—at any rate, certainly not less!

As far as I can see, there is simply nothing plausible for the advocate of the ratio view to say at this point. I believe that cases like this last one sound the death knell for the ratio view. It simply must be abandoned.⁶

5. Before turning to my own favored alternative principle, I want to consider the general, abstract question of how “optimistic” we should be with regard

⁵ Given the stipulation that A's peak is negative, this example might be rejected altogether by those who reject the retributivist claim that someone sufficiently vicious can deserve a level of well-being to the left of the origin. But even if you believe that no one deserves to suffer *overall*, so long as you believe that some deserve to be *punished*, corresponding difficulties for the ratio view will still arise. (Let the X axis represent not the overall level of well-being, but instead the magnitude of the reward or punishment that someone receives. If A deserves a punishment of -10, but can only be given one of -5, while B deserves a reward of +20, what size reward for B will satisfy comparative desert according to the ratio view?) I am indebted to Tom Hurka for discussion of this point.

⁶ I should note, in this regard, that since A's peak is negative and his level of well-being remains negative, and B's peak is positive and her level of well-being remains positive, this is not a case where *reversing* the ratio is called for. But even if it were, this would still give an unacceptable answer. Indeed, if instead of holding the relevant ratio to be -1 to 2 (A to B), we somehow claim that in this case it should be reversed to 2 to -1 (A to B), then we will have to say that if A is fixed at -5, B should be placed at +2.5. But this is an even more egregious violation of the intuition noted in the text—that if the less virtuous individual is given more than his peak, comparative desert cannot be satisfied by pushing the more virtuous individual *below* her peak.

to the possibility of satisfying comparative desert. There is a range of possible positions here.

The most pessimistic position holds that in certain cases it will be simply impossible to perfectly satisfy comparative desert, even if we have complete freedom with regard to the levels of well-being assigned to the relevant individuals. That is, according to the pessimist, in some cases no matter how we adjust the well-being of the various individuals, it will be impossible to fully satisfy the demands of comparative justice.

How could this be? To mention just one sort of suggestion, perhaps in some cases there is a kind of "many body" problem—where satisfying the demands of comparative desert between A and B, and between B and C, makes it impossible to simultaneously satisfy the demands of comparative desert with regard to A and C.⁷ (Fixing C's relation to B, and B's relation to A, automatically fixes C's relation to A; but it may do this in a way that doesn't meet the demands of comparative desert for A and C).

Although this is a possibility worth taking seriously, I am not inclined to accept it. I believe, instead, that optimism is appropriate: the demands of comparative desert can always be perfectly satisfied—at least so long as we have complete freedom with regard to the levels of well-being assigned to the relevant individuals.

This endorsement of optimism should not come as a surprise, given my earlier endorsement of the conjecture that when everyone is at their peak comparative desert is perfectly satisfied. Given complete freedom to locate individuals at whatever level of well-being we see fit, we can obviously place each person at their respective peak; and if the conjecture is correct this guarantees that comparative desert is perfectly satisfied. Thus anyone who accepts the conjecture—as I do—must embrace optimism as well. (Of course, one could accept optimism even without accepting the conjecture—so long as there is always *some* assignment of well-being that will satisfy comparative desert.)

Note, however, that all that is guaranteed by optimism is the ability to satisfy comparative desert given *complete* freedom with regard to assigning the well-being of the relevant parties. It remains possible, even if we are optimists, to insist that in at least some cases if there are some individuals whose level of well-being is *fixed*, it may be impossible to fully satisfy comparative desert (no matter how we adjust the well-being of the others).

Indeed, a moment's reflection makes it clear that we all must accept some kind of limitation to our optimism. For if enough people have their level of well-being fixed, in the "right" sort of ways, it will always be impossible to

⁷ George Sher tries to describe a case of this sort in *Desert* (Princeton: Princeton University Press, 1987), p. 33.

satisfy comparative desert. In point of fact, often all it takes is to have *two* people fixed in this way to make it impossible. Suppose, for example, that A and B are equally deserving, and so comparative desert demands that they have the same level of well-being. But obviously enough, if A's well-being is fixed at 10, while B's is fixed at 20, it will be impossible to satisfy this demand. Trivially, then, if two (or more) people are already fixed, it may well be impossible to perfectly satisfy comparative desert.

It seems, then, that there are only two kinds of optimists that we need to take seriously. One kind, let us call her a *limited* optimist, holds that if one has *complete* freedom with regard to assigning levels of well-being for the relevant parties, comparative desert can always be satisfied, but insists nonetheless that if even *one* person already has their level of well-being fixed it may prove impossible to fully satisfy comparative desert. In contrast, what we might call a *complete* optimist insists that even if there is already one person with their well-being fixed, so long as we have freedom with regard to assigning the well-being of others, it will still always be possible to satisfy comparative desert. (We can safely call this view "complete" optimism, since we have just seen that any view more optimistic than this—insisting that comparative desert can always be satisfied, even when the well-being of two individuals has already been fixed—is trivially false.)

However, once we are armed with these distinctions—between pessimism and optimism, and between the two kinds of optimism—it might be objected that my arguments in the previous section simply begged the question against the ratio view. In effect I simply assumed the truth of complete optimism, and then complained when the ratio view was not able to plausibly tell us how to satisfy comparative desert (in a manner compatible with the basic idea of ratio views). For example, in my final case, I fixed A's location at -5, and then complained that the ratio view gave an unacceptable answer when applied to that case. But this simply assumed that it was still possible to satisfy comparative desert in that case—despite the fact that one person's level of well-being was fixed. I thus implicitly presupposed the truth of complete optimism.

Admittedly, at the time we did not imagine the advocate of the ratio view complaining about my having set the case up in this way. It did seem that one could apply the ratio view even to a case of that sort, and when one did it gave an unacceptable answer. But now we must admit that there is a rather different type of answer available to those who are attracted to the ratio view. They could simply insist, in the kind of case that I described, that the ratio view is not to be "applied" there at all! They could insist that once the case has been set up in the way I described, comparative desert could no longer be perfectly satisfied. In short, perhaps the problem lies not with the ratio view but with the assumption of complete optimism.

A similar response is available to those who want to resist the suggestion, offered in reaction to the first example, that if we are to preserve the ratio view

we must sometimes reverse the ratio. For this case also involved fixing the level of well-being of one of the relevant individuals, and so perhaps advocates of the ratio view should simply have insisted that once this was done in the right way (once someone with a positive peak was fixed at a negative level of well-being) comparative desert could simply no longer be satisfied. Perhaps there was nothing wrong with the original ratio view at all, but only with the assumption of complete optimism.

Similarly, rejection of one or another form of optimism might be thought to suffice to disarm the various cases involving zero. Perhaps, for example, advocates of the ratio view should simply have insisted that if A's peak is 0, the relevant ratio cannot be defined. I complained at the time that saying this was incompatible with accepting the conjecture that comparative desert is indeed satisfied when everyone is at their peak. But we have now seen that the conjecture itself entails optimism (of at least a limited sort). So perhaps it is optimism that should be rejected (and the conjecture along with it), rather than the ratio view.

In light of this worry—that my arguments begged the question by presupposing one or another form of optimism—I must say at least a few words on behalf of optimism. One part of this defense I have, in effect, already given. For it does seem to me that it is extremely plausible to hold that when everyone is at their peak (that is, when noncomparative desert is perfectly satisfied) comparative desert is perfectly satisfied as well. This conjecture strikes me as extremely intuitive, and since it entails optimism, it seems to me that we should accept optimism as well.

Of course the conjecture only entails limited optimism, not complete optimism. Should we then accept complete, or only limited, optimism?

I believe that a further distinction is helpful at this point. I think it is useful to distinguish between cases involving only two individuals, and cases involving groups (three or more). For in cases involving only two individuals, I think the plausibility of complete optimism is quite high. If we are going to have a theory of comparative desert at all, and we are only dealing with a pair, then it seems to me that in principle it should be possible to perfectly satisfy the demands of comparative desert even when the level of well-being of one person is fixed. The basic idea of comparative desert, after all, is that my level of well-being should stand in the right relation to yours, in light of how our levels of virtue compare. But even if your own level of well-being is fixed, it seems to me, so long as we can assign me any level of well-being at all, it ought to be possible to get my level of well-being in the right relation to yours. Thus, so long as we restrict our attention to cases involving only pairs, complete optimism about comparative desert seems to me extremely attractive.

In contrast, it is far less clear that complete optimism is appropriate when we are dealing with groups of three or more. Indeed, as the remark about the

possibility of a “many body” problem suggests, it is not altogether clear that we should be optimists at all when it comes to cases involving groups.

Perhaps, then, we should be complete optimists about pairs, and pessimists about groups.⁸ In point of fact, I don't think this is quite right, since I do accept the conjecture, and I think it is plausible with regard to groups too, and not only to pairs. Thus I think we should accept at least limited optimism even with regard to groups. And perhaps we should go so far as to accept complete optimism with regard to groups as well.

Luckily, however, for present purposes I don't need to take a stand with regard to groups at all. Perhaps pessimism about groups is well taken, despite what I believe. Even if that is true, however, it remains the case that when it comes to pairs—when it comes to cases involving only two individuals—complete optimism remains an extremely attractive position.

It is important to note, therefore, that in my arguments against the ratio view, although I did indeed presuppose optimism, and sometimes presupposed complete optimism, it was always optimism with regard to pairs that was assumed. All of my cases involved only two individuals—A and B. If I am right in thinking that at least as far as pairs are concerned, complete optimism is a legitimate assumption (regardless of what we say about groups), then the arguments against the ratio view are successful after all.

6. Let me, therefore, propose an alternative comparative principle, which I will call the *Y gap view*.⁹ (The meaning of this name won't emerge until later, when we fill in some important details.) The guiding idea behind this alternative approach is this: comparative desert is perfectly satisfied when (and only when) the offense against noncomparative desert is the *same* for all relevant individuals.

Here's what I have in mind. From the standpoint of *noncomparative* desert what is important is that each person have what they absolutely deserve. But comparative desert, in contrast, is not primarily concerned with whether people have what they absolutely deserve. Rather, it is concerned with how people *compare* in this regard. That is to say, comparative desert is concerned with how people compare in terms of how they are doing relative to what they (noncomparatively) deserve.

Suppose, then, for example, that someone has less than they absolutely deserve. Clearly, this situation falls short of what is ideal from the noncomparative standpoint, and so it constitutes an “offense” against noncomparative desert. But so long as others similarly fall short, and fall short enough—so

⁸ It is interesting to note, in this regard, that Sher's pessimistic example (see n. 7) involves a group of three individuals. As far as I know he never argues for pessimism with regard to pairs.

⁹ It was suggested to me, independently, by both Alastair Norcross and Kyle Stanford. Clearly, I owe a tremendous debt to both of them.

long as the situation of others involves a similar offense against noncomparative desert—then no one has an advantage compared to the others with regard to how they are doing relative to what each absolutely deserves. The offenses against noncomparative desert are the *same*, and so *comparative* desert is satisfied.

Similarly, if someone has more than they absolutely deserve this constitutes an offense against noncomparative desert. But so long as the situation of others is such as to involve precisely the *same* offense against noncomparative desert, then no one has an advantage compared to any of the others with regard to how each is doing relative to what each absolutely deserves. Thus comparative desert will again be satisfied.

In short, what is important from the standpoint of comparative desert is not whether people are getting what they absolutely deserve, or offending against noncomparative desert—but only whether the offense against noncomparative desert, if any, is exactly the same for everyone. Even if my situation and your situation both involve some shortcoming from the perspective of noncomparative desert, so long as both of our situations involve the *same* offense against noncomparative desert, comparative desert will be perfectly satisfied.

This is, as I say, the guiding idea behind the Y gap approach, and I think it is an attractive one. Obviously enough, however, if we are going to take this general idea about comparative desert and turn it into a specific principle, we are going to have to specify precisely what it is for offenses against noncomparative desert to be the *same*.

I'll have a fair amount to say about this topic in a moment. But even before turning to it, I should note there is actually one kind of case that we are already in a position to discuss, namely, the case where there is *no* offense against noncomparative desert.

Suppose, for example, that you are at your peak. Then your situation is ideal from the standpoint of noncomparative desert; there is no "offense" against noncomparative desert at all. Accordingly, comparative desert will be satisfied just when my own level of well-being is also such as to perfectly satisfy noncomparative desert. Thus, comparative desert demands that I be at my peak as well. If, indeed, we are both at our respective peaks, then neither of our situations offends against noncomparative desert and so, trivially, the offense is the same for both of us (that is to say, none). Comparative desert is satisfied.

The first thing to notice, then, is that given the truth of the guiding idea behind the Y gap view, the truth of our conjecture falls out immediately: if we are both at our peaks, comparative desert is satisfied as well. Indeed, if you are at your peak, then the *only* way to perfectly satisfy comparative desert is to have me at my peak as well. Otherwise my situation will offend against noncomparative desert while yours does not, and so the offense against

noncomparative desert won't be the same for both of us. (The Y gap view thus also accommodates another attractive intuition noted previously.)

In short, if anyone at all is at their peak, comparative desert will be perfectly satisfied if and only if everyone is at their peak. This much we can safely say, even without having an account of what it is for offenses against noncomparative desert to be the same.

Other cases, however, will be less straightforward. If someone's level of well-being is fixed at a level above or below their peak, then an offense against noncomparative desert is inevitable. Despite this, of course, *comparative* desert can still be perfectly satisfied, so long as the situation of others is such as to involve precisely the *same* offense against noncomparative desert. But we don't yet know how to recognize when offenses are the same. Clearly, then, if we are going to say anything *determinate* about the demands of comparative desert in cases of this kind, we are going to need an account of the idea of a similar offense.

7. It seems to me that there are actually two conditions that must be met if the situations of two individuals are to offend against noncomparative desert in exactly the same way.

The first condition involves what we might think of as a *qualitative* constraint: similar offenses must be of the same *kind*. That is, if one person has more than their peak, the other must have more than their peak as well; and if one has less, then the other must also. Otherwise, even if the situations of both individuals involve offenses against noncomparative desert, the offenses cannot possibly be the same. It is, after all, one thing to be doing better than one absolutely deserves, quite another to be doing worse. Thus dislocation with regard to peaks must be in the same *direction*. We can call this the *symmetry principle*. It tells us that when either party is moved off their peak, comparative desert is only satisfied if both are moved in the same direction.¹⁰

By way of example, imagine that A is a less deserving individual than B, with a peak considerably further to the west. But let us suppose, as well, that A has more than his peak. Indeed, imagine that A's level of well-being is fixed at a particular point, well beyond what he absolutely deserves. (See Figure 4.2.) What then does comparative desert require with regard to B?

Given the symmetry principle, of course, we know that comparative desert will only be satisfied if B has more than her peak as well. Put in slightly

¹⁰ Strictly speaking, I suppose, we should distinguish between two claims. The first holds that symmetrical dislocation from the peak is required for similar offenses; the second holds that symmetrical dislocation is required to satisfy comparative desert. It is only when we add the guiding thought that comparative desert requires similar offenses that the second follows from the first. But since the present approach is indeed being guided by this thought, it would needlessly complicate the discussion to distinguish the two in what follows.




Fig. 4.2. Two individual desert graphs, with A's location fixed.

different terms, since A is on the eastern slope of his mountain, B should be on the eastern slope of her mountain as well.

But the symmetry principle—limited as it is to this qualitative requirement—cannot tell us *where* B should be located on her eastern slope. Symmetry tells us that B should be placed beyond her peak, but it does not tell us how far. We might say that it addresses quality but not quantity.

That is why the symmetry principle cannot, by itself, constitute an adequate account of comparative desert. Comparative desert requires that offenses against noncomparative desert be the same. But if B's situation is to involve an offense against noncomparative desert that is exactly the same as A's, it won't suffice that it be of the same *kind*; it must also be of the same *size*.

How then are we to measure the size of an offense against noncomparative desert? Happily, there is a very plausible proposal that suggests itself here. We can start by recalling that from the standpoint of noncomparative desert, the ideal situation is one in which the given individual is at their peak. In contrast, then, when someone is not at their peak—when they have less, or more, than they absolutely deserve—there is a drop off in value. That is why points on the individual desert graph to the west or to the east of the peak are at a lower level along the Y axis than the peak. Indeed, the further someone is from their peak, the greater the offense against noncomparative desert, and the lower the Y coordinate for the corresponding level of well-being. (Thus the individual desert graph takes on the shape of a mountain, with eastern and western sides sloping ever further down on either side of the peak.)

Thus we can measure the size of the offense against comparative desert by seeing how great the drop off is along the Y axis (as compared to the level along the Y axis that would obtain were the person actually at their peak). When someone is not at their peak there is a “gap” between the amount of




Fig. 4.3. The Y gap view.

intrinsic value (from the standpoint of noncomparative desert) that would obtain if they were at their peak, and the amount of intrinsic value that does in fact obtain, given their actual level of well-being. The greater the gap—the greater the drop along the Y axis—the greater the offense against noncomparative desert.

We thus arrive at a second condition on similar offenses. Similar offenses must involve a drop down the Y axis (as measured against the Y coordinate of the given person's peak) of precisely the same *size*. We can call this the *Y gap constraint* (hence the name for the overall account of comparative desert that incorporates it). It holds that comparative desert is satisfied only when the situation of each person is such as to involve a drop along the Y axis of exactly the same size.¹¹

Given the Y gap constraint we are now in a position to settle the question of where B should be located on her eastern slope. We have, of course, already stipulated that A's location on his eastern slope is fixed. Since he has more than he deserves this generates a drop off along the Y axis, a Y gap of a particular size. (See Figure 4.3.) But if B too is located at a point beyond her peak, this will also create a drop off in value—one that is greater in size the further she is beyond her peak. Imagine then that in point of fact B is at precisely that level where the extra, undeserved well-being (undeserved, that is, from the standpoint of noncomparative desert) yields a drop off in value of exactly the same size as the one created by A's having more than his peak. At this point

¹¹ Again (compare n. 10) we should, strictly, distinguish two claims. The first holds that Y gaps of the same size are required for similar offenses; the second holds that Y gaps of the same size are required to satisfy comparative desert. The second only follows from the first given the further thought that comparative desert requires similar offenses. But given our acceptance of this further thought, there is no need for us to further complicate the discussion by carefully maintaining the distinction between these two claims.

the Y gaps are exactly the same size, and the two situations offend against noncomparative desert by precisely the same amount. At this point, then, comparative desert is perfectly satisfied.

I suggest, then, that when the Y gap constraint is combined with the symmetry principle we have an adequate account of what it is for offenses to be exactly the same. They must be of the same kind, and they must be of the same size. Taken together, then, they provide what I take to be a highly plausible account of comparative desert: comparative desert is perfectly satisfied when offenses against noncomparative desert are the same size and the same kind. I call this account of comparative desert the *Y gap view*.

8. It is important to emphasize the point that the Y gap view makes essential use of the symmetry principle; it does not appeal to the Y gap constraint alone. For the Y gap constraint by itself would not in fact provide an adequate account of comparative desert. Just as the symmetry principle's concern for quality—when taken alone—is insufficient to determine the demands of comparative desert, so the Y gap constraint's concern for quantity—when taken alone—is insufficient as well.

The problem is this. The Y gap constraint tells us how far any given individual should be from their peak, given the location of another. In particular, it insists that the situation of each must involve a drop down the Y axis of precisely the same size. But this constraint can typically be satisfied in *two* ways, since a gap of a given size can normally be produced by locating a person on either their eastern slope or their western slope. (The sole exception, of course, is when one person is at their peak. Then there is no gap, and the other must be at their peak as well.)

For example, imagine once again that A's location is fixed on his eastern slope, generating a Y gap of a particular size. Obviously, then, one way to satisfy the Y gap constraint will be to locate B appropriately on her eastern slope, specifically at B₁. (See Figure 4.4.) But satisfying the Y gap constraint does not actually require that B be located on her eastern slope at all. For even if B is placed on her *western* slope, this will still generate a drop down the Y axis. Thus, if we imagine B located at B₂—where the size of the Y gap is also precisely that of the gap produced by A's location—the Y gap constraint will still be satisfied. Either of the two points marked out on the sides of B's mountain will satisfy the Y gap constraint.

Obviously, however, it is quite implausible to suggest that if B is placed at B₂ this satisfies comparative desert. A is considerably less virtuous than B, and has far more than he absolutely deserves. Clearly, in a situation like this, comparative desert cannot be satisfied with B having *less* than she deserves.

Luckily, this is not any kind of objection to the thought that comparative desert is satisfied when offenses against noncomparative desert are exactly the same. For here, even though A's situation and B's situation offend against




Fig. 4.4. The inadequacy of the Y gap constraint taken alone.

noncomparative desert by the same amount, they are not at all the same. They are similar in size, but not in kind. What we learn, then, by thinking about a case like this is simply that comparative desert requires more than that offenses be of the same size.

That is why the Y gap constraint, by itself, does not provide an adequate account of the demands of comparative desert. We must supplement it with the symmetry principle. (It is the symmetry principle that rules out our locating B at B₂. By hypothesis, A has more than his peak, and so B must end up with more than her peak as well.)

Thus neither the Y gap constraint nor the symmetry principle *taken alone* constitutes an adequate account of comparative desert. But both, I think, make plausible claims about what comparative desert requires. And taken together they do in fact yield a plausible proposal. Taken together they yield the Y gap view, according to which comparative desert is perfectly satisfied precisely when the situations of the relevant people involve similar offenses against noncomparative desert—offenses of the same size and of the same kind.

(Although the particular example we have been discussing has involved someone fixed at a level higher than their peak, I trust it is clear how the Y gap view would deal with a case where someone has *less* than their peak. If you, say, have a level of well-being that is less than what you absolutely deserve then comparative desert will demand that I too have less than I absolutely deserve—just enough less, in fact, so that the Y gap created by my own situation is exactly the same size as the Y gap created by your situation. At this point our two situations will offend against noncomparative desert in precisely the same way, and so *comparative* desert will be perfectly satisfied.)

9. In section 4 I argued for the rejection of the ratio view on the ground that it faced a series of problems in cases involving negative numbers and zero. In contrast, it is worth noting, the Y gap view has no problems handling cases of

this sort at all. To support this claim, let me quickly work through the same three examples that were so problematic for the ratio view, to see how readily the Y gap view deals with them.

In the first case, you will recall, we stipulated that A had a peak of 10, and B a peak of 20, but that A's well-being was fixed at -10. The problem for the ratio view, of course, was that unless we modified our understanding of that view (by reversing the relevant ratio in such cases) it implausibly entailed that B should end up worse than A, despite being more absolutely deserving. But the Y gap view isn't similarly threatened. Given symmetry, of course, it is certainly true that since A has less than his peak, B should have less than her peak as well. But the Y gap constraint requires that B's drop down the Y axis from her peak be the same as A's drop from his peak—and this has the implication that B must end up no further to the west of her peak than A is to the west of his.¹² Thus B will end up—appropriately enough—to the east of A, since B's peak is to the east of A's.

In the second case, A has a peak of 0, B has a peak of 10, and A is actually at his peak. Here the problem was that the ratio view either had to claim, implausibly, that comparative desert is satisfied regardless of where B is located (even if B is worse off than A!), or it had to claim, also implausibly, that comparative desert would not be satisfied even with B at her own peak. But the Y gap view avoids both of these unacceptable alternatives. Since A is at his peak, his situation involves no offense against noncomparative desert; and so the Y gap view appropriately tells us that comparative desert demands that B must be at her peak as well.

Finally, in the third case, A has a peak of -10, B has a peak of 20, and A is actually at -5. The ratio view absurdly implied that in the face of A's having more than he absolutely deserved, B should be given *less* than she absolutely deserved. And once again, the Y gap view easily avoids this absurdity: since A has more than he absolutely deserves, symmetry demands that the same should be true of B.

In sum: the ratio view faced a series of problems when it confronted cases involving zeros and negative numbers. In contrast, such cases pose no particular difficulties for the Y gap view at all.

¹² Where, exactly, should B be located, according to the Y gap view? That will depend on one of the complications in the theory of noncomparative desert that we noted at the end of section 1. I argued there for the existence of what I called "bell motion," and if we accept it then the western slope of B's mountain will actually be steeper than the western slope of A's mountain. This means that when it is a matter of having less than what each absolutely deserves, it takes a smaller shortfall in well-being for B to have a Y gap comparable to A's (since each unit change to the west along the X axis results in a larger change along the Y axis). So, given bell motion, comparative desert will actually demand that B be placed somewhat closer to her peak than A is to his. How much closer, though, will depend on details about the *rate* of bell motion that we cannot pursue here. What is crucial for our present purposes, however, is just this. Unless one accepts *reverse* bell motion—the quite implausible view that the western slopes are *steeper* for the less virtuous—the Y gap view will never place B at a lower level of well-being than A.

This does not show, of course, that the Y gap view avoids all such difficulties, and that its answers are plausible—and free from absurdity—across the board. It would take too long to run through the full range of possible cases, nor is there space here to consider the various other objections that might be raised against it. But I believe that the requisite, fuller investigation would simply confirm what we already now have good reason to believe—that the Y gap view provides the basis for a plausible general account of comparative desert.¹³

10. It may be helpful, however, to say a bit more about the Y gap constraint. As the name itself reminds us, I've chosen to state it in terms of *Y gaps*—the distance one lies down the Y axis from one's peak. In the terms of this language, for example, the constraint requires that the Y gaps be the same for both A and B. But it might seem that this same requirement could have been expressed in a slightly different—and a much simpler—way, namely, that A and B must have the same location on the Y axis.

The thought here is fairly straightforward, and is especially easy to grasp if one looks again at Figure 4.3. A is considerably less virtuous than B, and so has a peak located considerably further to the west. But although A and B differ in terms of the X coordinates of their peaks, the Y coordinates of the two peaks are the *same*: if each gets what they deserve, this is equally good from the standpoint of desert. Now the Y gap constraint requires that the distance down the Y axis from their peaks be the same for both A and B. But if both "start" at the same point along the Y axis (since the peaks have the same Y coordinates) and both "come down" the same distance along the Y axis (so as to satisfy the Y gap constraint), then they must both "end up" at the same point along the Y axis as well. In short, or so it seems, the Y gap constraint could be more easily stated as a requirement that both A and B be so located as to have the same Y coordinates (that is, they must be at the same level along the Y axis).

Now in point of fact this argument only succeeds if it is indeed true that all peaks have the same Y coordinate. Obviously enough, if some peaks are located at a higher level along the Y axis than others, then comparable drops down the Y axis will sometimes require *differing* locations along the Y axis. Thus the "simpler" formulation—where the relevant parties must end up with the same Y coordinates—will only truly be equivalent to the Y gap constraint if all peaks are at the same height.

But this, of course, is a question about the nature of the *skyline* (the line consisting solely of possible peaks); and as we noted in section 1 there is more

¹³ A "basis" rather than "a complete account," since there are further issues that we have not tried to address here—most pressingly, issues concerning the many cases where the demands of comparative desert are only *imperfectly* satisfied.

than one position which can be plausibly taken here. It is certainly true that one possible view holds that all peaks have the same Y coordinate (that is, that the skyline is a straight line parallel to the X axis). On this view, then, though people differ in terms of what they deserve, it is equally *important* that everyone get what they deserve. But this is not the only attractive possibility. As we previously noted, it might be, instead, that it is more important that some people—the more virtuous, or the more vicious—get what they deserve. On this alternative account, the skyline would have a V shape (with the bottom of the V located at the origin).

If the skyline has a V shape then the argument for the simpler formulation does not go through. Rather, if A's peak is lower than B's, satisfying the Y gap constraint requires that A's actual location be further down the Y axis than B's. (See Figure 4.5.)

So given a V-shaped skyline, the Y gap constraint gives rather different answers than those that would be given by a requirement that everyone have the same Y coordinates. Perhaps, however, forced to choose between the two, we should pick the latter rather than the former?

I don't believe this would be a plausible choice. For it would mean that even in the situation that is ideal from the standpoint of noncomparative desert—with both A and B at their respective peaks—comparative desert wouldn't be satisfied (since B's peak is higher along the Y axis than A's). And this, of course, violates the extremely plausible conjecture that when noncomparative desert is perfectly satisfied, comparative desert is perfectly satisfied as well.

At a minimum, then, those who accept a V-shaped skyline (as well as various other more exotic possibilities) must resist the proposal to restate the Y gap constraint in terms of a requirement that people end up with the same Y coordinates. They must retain the original formulation of this constraint, in terms of the drop down the Y axis.

But in point of fact even those who accept the suggestion that the skyline is a straight line parallel to the X axis have reason to retain the original




Fig. 4.5. The Y gap view with a V-shaped skyline.

formulation. To be sure, since all peaks have the same height on this view, the Y gap constraint will be satisfied just in case people's actual locations have the same Y coordinates. From a practical point of view, then—if this view about the skyline is correct—focusing on the Y coordinates may well be an easy way to determine whether the Y gap constraint has been satisfied. But for all that, there is still a philosophical advantage in having our “official” statement of the Y gap constraint given in terms of a requirement for a similar drop down the Y axis. For only this formulation wears its rationale on its sleeve.

After all, in directing us to compare the actual situation to the ideal one (that is, to the situation where the person is located at their peak), this formulation forces us to view the actual situation in terms of the extent to which it falls short of the demands of noncomparative desert. It thus expressly embodies the very attractive idea that comparative desert is a matter of looking for similar *offense* against noncomparative desert. More particularly still, it expresses the plausible idea that offenses against noncomparative desert must be the same *size* if comparative desert is to be satisfied.

There are, then, at least two reasons to prefer to state the Y gap constraint—and the Y gap view that incorporates it—in terms of the drop off along the Y axis. On the one hand, unlike the proposed alternative, it gives plausible answers regardless of the shape of the skyline; and, on the other hand, it is stated in such a way as to make plain its connection to a plausible overall conception of comparative desert.

11. It may also be helpful to contrast the Y gap constraint with a somewhat different proposal, one that we can call the *X gap constraint*. According to this view, the size of an offense against comparative desert should be measured, not in terms of Y gaps (the distance down the Y axis from the peak), but in terms of X gaps (the distance along the X axis from the peak). If your peak is 20, for example, but you are in fact at 25, then your situation involves an X gap of 5. Comparative desert will be satisfied, then, only if I too am at a level of well-being 5 units away from my peak.

Of course, like the Y gap constraint, the X gap constraint can typically be satisfied in two ways—by giving me either the appropriate amount too much, or the same amount too little. In the case just considered, for example, if my own peak is 10, then my situation will involve an X gap of 5 regardless of whether I am located at 5 units of well-being or 15. Put in slightly different terms, the X gap constraint is typically indifferent between placing someone on their western slope or on their eastern slope. (As before, the sole exception is when someone is at their peak. Then the X gap is zero, and so others must be placed at their peaks as well.)

But like the friends of the Y gap constraint, advocates of the X gap constraint can appeal to the symmetry principle to deal with this problem. Given symmetry, since you have more than your peak, I must have more than

my peak as well; otherwise the offense against noncomparative desert will not be of the same kind. That settles the question of whether to place me on my eastern slope or on my western slope. But now, instead of appealing to the Y gap constraint to tell us how *much* more than my peak I should have, we appeal to the X gap constraint. Your X gap is 5, and my X gap must be the same size as yours; so I should be placed 5 points beyond my peak, at 15.

Thus we arrive at the X gap *view*, according to which comparative desert demands similar offenses, where offenses must be of the same kind and the same size—but size is measured in terms of X gaps rather than Y gaps.

As it turns out, this new view can be restated more simply: if one person absolutely deserves a certain amount more than another, then comparative desert demands that that person have that much more than the other.¹⁴ (And in the limit case, where both are equally deserving, both must be at the same level of well-being.) No doubt, the X gap view is more familiar in this alternative formulation—and indeed in section 3 I mentioned it, in similar terms, as one possible view.

Of course, if we do restate the X gap view along these simpler lines, we run the risk, once again, of losing sight of the underlying rationale behind the approach—the guiding thought that comparative desert demands similar offense. But for the moment we can let this point pass.

Apparently, then, we face a choice between two alternative accounts of comparative desert, the X gap view and the Y gap view. Which should we prefer?

It might seem, however, that there is in fact no need to choose at all. For it is tempting to believe that the two views are actually equivalent. And once again a look at Figure 4.3 should make the thought here easy to understand. A is considerably less virtuous than B, and so his peak is located further to the west. Indeed his entire mountain is, in effect, shifted to the west of B's. But despite this difference in location, the mountains are otherwise similar. In particular, the slope of the eastern side of A's mountain is the same as the slope of the eastern side of B's mountain. But this means that the Y gaps for A and B can be the same size if and only if the X gaps are the same size as well. (If the two eastern slopes are the same, then similar dislocations along the X axis from the peak must produce similar dislocations along the Y axis.) Thus it doesn't actually matter whether we measure the size of the offense against noncomparative desert in terms of X gaps or Y gaps, so long as we are consistent in our approach; the two approaches must always agree as to whether or not we are dealing with offenses of the same size. (I've given the argument

¹⁴ Suppose that B's peak is N units to the east of A's. According to the X gap view, comparative desert demands that dislocation along the X axis from one's peak must be in the same direction, and by the same amount. Thus B must *end up* at a point N units to the east of A as well, if comparative desert is to be satisfied. (And the same is true, obviously, even if $N = 0$.)

with regard to eastern slopes; but a similar argument could obviously be mounted for western slopes as well.)

Clearly, this argument only succeeds if the eastern sides of the two mountains do indeed have the same slope. But this is, in fact, a controversial claim, and it brings us back to one of the complications that I mentioned at the end of section 1. I noted there the plausibility of accepting bell motion, according to which the mountain swings like a bell (if we think of the peak as fixed), to the left for vicious individuals and to the right for virtuous individuals. In particular, I think it plausible to hold that the slope of the eastern side of the mountain grows gentler with more virtuous individuals, so that if one must overcompensate either a more virtuous individual or a less virtuous individual by a certain amount, it is less bad to overcompensate the more virtuous individual.

But if we accept this idea, then we must abandon the claim that Y gaps will be the same if and only if X gaps are the same. In Figure 4.6 I've drawn a graph that, unlike our earlier graphs, displays bell motion and so makes this clear. (I've exaggerated the rate of bell motion to make it easy to see; but this won't affect the relevant points.) With A's actual location fixed at the indicated point on his eastern slope, we now face the question of where to place B on her eastern slope. But since B is more virtuous than A, her eastern slope is gentler, and thus it takes a greater dislocation along the X axis to produce a comparable dislocation along the Y axis. Thus the X gap view and the Y gap view differ concerning the demands of comparative desert. The X gap view will claim that B should be placed at B_2 (where B's X gap is the same size as A's, but the Y gap is much smaller). And the Y gap view will claim that B should be placed at B_1 (where B's Y gap is the same size as A's, but the X gap is much greater).




Fig. 4.6. The difference between the Y gap view and the X gap view.

In short, given bell motion, the Y gap and the X gap views come apart. Which, then, should we prefer?

Conceivably, we could try to settle this question by appeal to intuitions about particular cases. We could construct examples, determine the precise answers given by the X gap view and the Y gap view, and see which struck us as more plausible. But such an approach would face at least two problems. First, until we develop the account of bell motion, we cannot say precisely how the slopes of the different desert graphs will vary (we cannot assign particular slopes), and this means, unfortunately, that we are not yet in a position to derive determinate answers to compare the two views. (If we stipulate the size of A's X gap, for example, we aren't yet in a position to determine the size of his Y gap; and if we stipulate his Y gap, we can't determine his X gap.) Second, even if we were able to produce determinate numbers for particular examples, it is not clear to me that we have sufficiently finely attuned intuitions to allow us to have any confidence in our preferences for one set of answers here rather than another.

Presumably, however, this still leaves us with the possibility of evaluating the alternative rationales that can be offered on behalf of the two views. Case specific intuitions aside, then, what general philosophical reasons do we have to prefer one approach over the other?

Typically, of course, the X gap view is presented without explicit appeal to any underlying rationale at all. But I have suggested, in effect, that the X gap view is most charitably understood as being based on the same guiding thought as the Y gap view, namely, that comparative desert demands similar offenses against noncomparative desert. The two views differ only in terms of how the size of the offense is to be measured. So our choice between the two views boils down to this: which is the relevant measure of the *size* of an offense against noncomparative desert—the X gap or the Y gap?

To be sure, there is a sense in which both of these measure *something* objective about the offense. When I have more or less than my peak, there is a fact of the matter concerning how much more or less I have; and this is measured by the X gap. And similarly, there is a fact of the matter concerning how much worse this state of affairs is than the one that would obtain if I were, instead, at my peak; and this is measured by the Y gap. So both have a claim to being genuine measures.

But our concern here is with the question of which measure is the *relevant* one for determining whether comparative desert is satisfied. And from this perspective, I believe, there is clear reason to prefer the Y gap.

For what we want to know is this: if the situations of two people both offend against noncomparative desert, are the situations equally *offensive*? Obviously enough, two situations can be similar in any number of ways and yet differ in terms of how bad they are from the standpoint of desert. (For example, suppose you and I have the same level of well-being; but I am at my

peak, while you are below yours.) What we want to know, then, is how much of a drop off in value there is from the standpoint of noncomparative desert.

But this is precisely what the Y gap is stipulated to measure. For the Y axis represents goodness from the standpoint of (noncomparative) desert. Therefore, gaps along the Y axis (relative to the peak) are precisely the measure of the extent to which a given state of affairs offends against noncomparative desert. It is thus precisely the Y gap to which we should appeal, in deciding whether or not two situations are equally offensive from the point of view of noncomparative desert.

In contrast, there is no particular reason to think that the X gap similarly measures the extent to which a given state of affairs offends against noncomparative desert. On the contrary, I suspect that if we are tempted to appeal to X gaps this is because we unwittingly assume that measuring in terms of them will be equivalent to measuring in terms of Y gaps. (Normally, that is, we simply overlook the existence of bell motion, implicitly assuming that differing mountains have different peaks but the same slopes, and then implicitly recognize that, given this assumption, X gaps will be the same if and only if the Y gaps are the same.) But once we bear firmly in mind that X gaps and Y gaps can come apart—and will often do so, given bell motion—we should recognize that for the purposes of determining the size of the offense against noncomparative desert, it is the Y gap that is relevant, and not the X gap. Indeed, there is no particular reason to be concerned with X gaps, *per se*, at all.

In short, forced to choose between the Y gap view and X gap view, we have reason to believe it is the former, rather than the latter, which provides the best account of comparative desert.

Of course, not everyone accepts the existence of bell motion, and so not everyone will agree that the X gap view and the Y gap view can come apart in this way.¹⁵ As we have already noted, if the slopes of different mountains are actually the same, then indeed Y gaps will be the same if and only if X gaps are the same. From a practical point of view, then, it won't be pressing for us to choose between the two views.

Nonetheless, as we now realize, even here there is some philosophical advantage, at the very least, in having the official account of comparative desert stated in terms of Y gaps rather than X gaps. For only talk of Y gaps explicitly directs our attention to what is truly of direct relevance—the extent to which a given situation offends against noncomparative desert. In short, even if the two views are otherwise equivalent (even if we reject bell motion), the Y gap view remains more philosophically perspicuous, and is thus to be preferred.

¹⁵ Strictly, of course, bell motion is not the only account according to which the slopes of mountains can differ from individual to individual (though I believe it is the most plausible account). Thus it is not the only account according to which the Y gap view and the X gap view will diverge.

We thus have reason to prefer the Y gap view regardless of whether or not we accept bell motion. For comparative desert demands similar offenses against noncomparative desert—offenses of the same kind and of the same size: Measuring the size of any given offense in terms of Y gaps makes it clear just how our approach fits in as part of this plausible overall conception of comparative desert.