A Solution to the Problem of Indeterminate Desert[*]

Brad Skow

May 14, 2009

Abstract

A *desert-sensitive moral theory* says that we have some obligation to ensure that people get what they deserve, or are treated as they deserve to be treated. Some popular forms of consequentialism are desert-sensitive. But where do facts about what people deserve come from? If someone deserves a raise, or a kiss, in virtue of what does he deserve those things? One plausible answer is that what we deserve depends, at least in part, on how well we meet our moral requirements. The wicked deserve to suffer and the decent do not. Shelly Kagan [2006] has argued that this plausible answer is wrong. But he is mistaken. I will show how to formulate a desert-sensitive moral theory (and also a desert-sensitive form of consequentialism) on which this answer is correct.

1 A Problem for Consequentialists

Desert plays an important role in our moral thought. Something has gone off, we think, when wicked people prosper and decent people suffer. And that is because the wicked do not deserve to prosper, and the decent do not deserve to suffer.

[*]It is Dan Greco's fault that I wrote this paper. I thank him, Caspar Hare, and Agustin Rayo for discussing this material with me. I would also like to thank Fred Feldman for checking the proof in the appendix.

If these claims about desert are correct, then they constitute a problem for simple versions of consequentialism. Roughly speaking, consequentialists say that an action is right if no alternative to it leads to an outcome with higher value.[1] For this to make sense consequentialists need a way to assign values to outcomes, or worlds. Suppose that a consequentialist says that the value of a world is just equal to the sum of the values of the lives of the people in that world. That is, to determine the value of a world, we find the number that represents how good each person's life is *for him* (his level of individual well-being) in that world, and add those numbers up. When a wicked person prospers, his life goes better for him; the value of his life for him goes up. So this version of consequentialism must say that it is a good thing when the wicked prosper: a world in which some wicked person prospers is better (other things being equal) than a world in which that person (is wicked and) suffers. And so this version of consequentialism must say that the prospering of the wicked is something that (other things being equal) we ought to bring about: if there are only two acts open to me, and the only difference in the outcomes the acts will produce is that one causes a wicked person to suffer, and the other causes that wicked person to prosper, then I ought to perform the latter action. Many consequentialists do not want to have to say this.

One way around this problem is to "adjust utility for desert." Change the way the value of a world is determined so that facts about what people deserve play a role in determining that value. One way to formulate consequentialism to take desert into account proceeds roughly as follows. Distinguish between two kinds of value someone's life has. There is the value his life has for him—his level of welfare, or individual well-being. And there is the value his life contributes to the value of the world. The simple version of consequentialism above said that these two values are always equal in magnitude. A desert-adjusted axiology denies this. It says, instead, that the value someone's life contributes to the value of the world is a function of the extent to which the value of that person's life for him matches up with the value of the life he deserves to lead. If

[1] I believe that a better and more precise formulation of consequentialism uses the concept of a world being accessible to an agent at a time. It says that it is right for S to perform a at t iff S performs a in (one of) the best worlds accessible to S at t. Feldman [1997c] defends this formulation and explains its advantages. For the purposes of this paper, though, the rougher formulation will do.

someone who deserves misery lives a life full of pleasure, then his life has positive value for him, but contributes negatively to the value of the world.[2]

This kind of desert-adjusted axiology allows consequentialists to say that the prospering of the wicked is bad—if it can be made to work. But there are reasons to think that it cannot be made to work. One difficulty is that it may leave the values of some worlds undefined. This is the problem of *indeterminacy.* Let us continue to assume that the decent deserve to prosper and the wicked deserve to suffer. That is, people who act morally deserve to live good lives full of pleasure, while people who do the wrong thing deserve to suffer. Now suppose that I am alone in the universe, and that I cause myself pleasure. Then it is consistent to suppose that I deserve the pleasure, and so the world has positive value. (For then, since I do the right thing, I am decent, and deserve the pleasure.) It is also consistent to suppose that I do not deserve the pleasure, and so the world has negative value. So this world has no determinate value. But consequentialism, as standardly formulated, requires a determinate assignment of value to every world.

Shelley Kagan [2006] has developed this line of thought and argued, in general, that it just is not true that what and how much you deserve depends, at least in part, on how well you live up to your moral requirements. His arguments are not aimed specifically at consequentialism. Instead, he argues that in any desert sensitive moral theory what one deserves is independent of whether one conforms to the requirements of morality.

Kagan does not present a general proof of this claim; instead he discusses many possible theories of the way that desert may depend on conformity to moral requirements, and demonstrates that they all suffer from indeterminacy. For all that he has said, there may be a theory that he has not considered that does avoid indeterminacy. I will defend the view that Kagan attacks by presenting such a theory—actually, two such theories. The first one I will present resembles one that Kagan considers; the main difference from Kagan's, and the reason why it succeeds, is that it is more mathematically sophisticated. But this theory is not, on its face, a version of consequentialism. I will conclude by suggesting a way for consequentialists to use a similiar strategy to reply to

[2]This idea is developed in more detail in [Feldman 1997a]. See also [Arrhenius 2006] for more discussion. The precise details will not matter for what I want to say.

Kagan's arguments, and make some remarks about the plausibility of that strategy.

2 Re-Stating the Problem

I began by presenting the problem of indeterminacy as a problem for versions of consequentialism that have a desert-sensitive axiology. It will be useful to present the problem of indeterminacy again in a more abstract framework. This framework does not assume that some form of consequentialism is true, and simplifies morality in other ways to make the problem easier to discuss. (This is the framework in which Kagan presents the problem.)

Here is the framework. First, there is only one moral requirement: to give others what they deserve. Someone is good (let us say) if he conforms to the requirements of morality, and bad if he does not. (This is a stipulative definition of "good." It may not match the meaning the word has in ordinary usage.) Second, there are only two things that people can deserve: love and hate. Each person must either love or hate each person (and cannot do both). Finally, desert depends on conformity to morality in the simplest way: someone deserves love if and only if they are good (and so deserves hate if and only if they are bad.)

Now we can describe a situation in which indeterminacy arises in this framework. Suppose that there are two people, Frank and Stella. Suppose that Frank and Stella love each other. Does this determine whether they are good or bad? The answer is "no." It is consistent to suppose that both are good. For if Frank is good, then he deserves love; Stella loves him, so gives him what he deserves, so Stella is good. Since she is good, she deserves love, which is what Frank gives her; so Frank is good. Consistent. But it is also consistent to suppose that both are bad. For if Frank is bad, then he does not deserve love; so Stella does not give him what he deserves, so Stella is bad. Since Frank loves Stella, he does not give her what she deserves, so he also is bad. This is also consistent. That is the problem of indeterminate desert.[3]

[3] I introduced the problem of indeterminacy for consequentialists with a scenario with just one person in it. Indeterminacy can also arise in a scenario with just one person in this more abstract setting. Suppose that the only moral requirement is to give each person what he deserves, and suppose that there is one person who loves himself. But it is easy to avoid this indeterminacy, by

It is worth noting that indeterminacy is not the only problem for this theory. Another problem is that it is *inconsistent.* Suppose that Frank loves Stella but Stella hates Frank. Then there is *no* consistent way to say what Frank and Stella deserve. (If Stella deserves love, then Frank is doing what he ought, so he deserves love, so Stella does not do what she ought, so Stella deserves hate. Similarly, if Stella deserves hate, we can show that she deserves love. This is an analogue of the version of the liar paradox in which sentence $L1$ says that $L2$ is true, and $L2$ says that $L1$ is false.) My solution to the problem of indeterminacy will also solve the problem of inconsistency.[4]

3 A Preliminary Solution

Let us start by thinking through the problem of indeterminacy again, in more abstract terms. We start with certain facts that are given: we are given whether Frank and Stella love each other, and we are given what the moral requirements are. Then we begin with some supposition about how well Frank and Stella conform to the moral requirements. From this supposition, we can compute what each of them deserve, and then from that compute how well they conform to the moral requirements. We have gone in a circle from a level of conformity to a level of conformity. The supposition we started with is consistent if the level of conformity we get after we have gone around the circle is the same as the level we supposed to obtain at the beginning. The problem is that there are many consistent suppositions.

Now I can briefly summarizing my strategy for solving this problem. The strategy has two components. First, I will expand the number of possible levels of conformity to morality that Frank

saying either that each person is good by "default," or that morality does not require each person to give himself what he deserves; morality only requires that we give *others* what they deserve. Kagan demonstrates that neither of these solutions works in scenarios with more than one person.

Of course, it is not obvious that *consequentialists* can use either of these strategies to avoid indeterminacy in one-person scenarios. (Though Ted Sider's [1993] "self/other utilitarianism" may be able to.) But the consequentialist solution I discuss below does not need them.

[4] In fact, the problem of inconsistency is much more serious than the problem of indeterminacy. The might be ways to live with a theory that permits indeterminacy; but (deviant logicians aside) there is no way to live with a theory that is inconsistent. I focus more attention on the problem of indeterminacy because Kagan focuses more attention on the problem of indeterminacy.

and Stella can have. This has the effect of *expanding* the number of "starting points" there are for going around the circle. Second, I will add to the number of moral requirements. This has the effect of changing what happens when you go around the circle, from supposed conformity levels to final conformity levels. The net effect will be to leave exactly one consistent supposition about Frank and Stella's levels of conformity.

Now I will explain my solution in detail. The problem of indeterminacy as presented presupposes that what you deserve and whether you conform to morality come in discrete quantities. In the simple example each comes in only two: you either deserve love, or hate, and you either conform to the demands of morality, or you do not. (Kagan also discusses a theory in which there are more than two levels of desert and of degree of conformity—but still only finitely many.) The first part of my solution is to move to a theory that allows these quantities to vary continuously, and so to come in infinitely many levels.

It is plausible that allowing these quantities to vary continuously will solve the problem of inconsistency. (It looks to help with the liar paradox: just assign each sentence a degree of truth of 1/2.[5]) But it is not obvious that continuous variation will help with the problem of indeterminacy. Won't continuous variation just provide *more* opportunities for indeterminacy? The answer is "yes": continuous variation alone will not solve the problem. I will come to what else we need to add soon. First we need to figure out how to think about varying these quantities continuously.

Each person has a degree of conformity to the requirements of morality, which I will represent with numbers: 1 is perfectly saintly, someone who always does what he ought, and 0 is completely vile, someone who never gets it right. But I want to allow any degree of conformity between these. For scenarios with two people, like Frank and Stella, the joint state of conformity to moral requirements will be a pair of real numbers (f, s), each between 0 and 1. Let M be the set of "conformity states" $[0, 1] \times [0, 1]$.

Not just conformity to moral requirement, but also love comes in degrees. We need a convention for representing degrees of love. I will say that love of degree 1 is perfect love, while

[5] In fact, this approach to solving the liar paradox is no longer very popular. See [Field 2008: sections 4.4-4.5] for a summary of the problems it faces.

"love" of degree 0 is perfect hatred. Love of degree 1/2, then, is complete indifference. (I am using "love" in a technical sense here, but it is related to the ordinary sense. I love someone to some degree in the ordinary sense if the degree to which I love them in this technical sense is greater than 1/2. If we say that a number y is the opposite of x just in case $y = 1 - x$, then in this scheme love is the opposite of hate.[6])

Each person also has a degree to which he deserves to be loved. In light of what degrees of love can be, the set of "desert states" is easy to characterize: the set D of desert states is also $[0, 1] \times [0, 1]$. An ordered pair (x_1, x_2) in this space represents a state in which Frank deserves to be loved to degree x_1 and Stella deserves to be loved to degree x_2.

This is to be a theory that says that what people deserve depends, at least in part, on how well they conform to the requirements of morality. So somehow facts about how much love Frank and Stella deserve are determined by the facts about how well they conformity to morality. We can represent this dependence by a function h from M to D. So h works like this:

$$h(\text{degree of conformity}) = (\text{degree of love deserved}).$$

If, for example, Frank's degree of conformity is .7, and Stella's is .2; and if in virtue of this Frank deserves to be loved to degree .9, and Stella to degree .5, then we write $h(.7, .2) = (.9, .5)$. What is the function h? The obvious choice is the identity function: the degree to which someone deserves to be loved is just equal to their degree of conformity. Those who always act rightly deserve perfect love, those who sometimes get it right deserve some love, and the absolutely vile who always do the wrong thing deserve complete hatred.

That is the story about how desert is determined by conformity to morality. As for conformity, we have only one moral requirement so far: an obligation to love others iff they deserve it. You satisfy this obligation perfectly when you love someone to just the degree to which they deserve to be loved. What about the case where you do not love someone to the appropriate degree?

[6] Even though, as Elie Wiesel says, the opposite of love is not hate, it's indifference.

One, simple, approach is to say that the degree to which you satisfy your obligation is determined by the "distance" between the degree to which you love them, and the degree of love they deserve. The closer these two degrees are, the closer you are to perfectly satisfying your obligation. So if you love that person to degree x, and they deserve love of degree y, then we could say that the degree to which you meet your obligation is just $1 - |x - y|$. Then, as we want, when both degrees match up, so $|x - y| = 0$, you satisfy your obligation to degree 1. When you are getting this completely wrong, hating those who deserve perfect love, so $|x - y| = 1$, you satisfy your obligation to degree 0.

Suppose it has already been determined who loves whom, and to what degree. Given these facts, the facts about how much each person deserves to be loved serve to determine each person's level of conformity to morality. We can represent this by a function g from D to M:

$$g(\text{degrees of love F and S deserve}) = (\text{how well each conforms to morality}).$$

If l_f is the degree to which Frank loves Stella and l_s is the degree to which Stella loves Frank, and if Frank deserves to be loved to degree x_1, and Stella deserves to be loved to degree x_2, then Frank's degree of conformity is $1 - |l_f - x_2|$. (The number x_2 appears because Frank's level of conformity depends on whether he loves Stella as much as she deserves to be loved, and it is x_2 that represents Stella's level of desert.) So the mathematical representation of g is

$$g(x_1, x_2) = (1 - |l_f - x_2|, 1 - |l_s - x_1|). \tag{1}$$

We can use these two functions g and h to build a function from M, the conformity states, to itself. The function simply takes a conformity state (f, g) to $g(h(f, g))$. Call this function "k": starting from a claim about Frank and Stella's level of conformity, this function takes us to what they deserve, and then (given how they love) back to their level of conformity. Abstractly, k looks like this:

$$k(\text{how well F and S conform to morality}) = (\text{how well F and S conform to morality}).$$

Since h is the identity function, the mathematical form for k is just

$$k(f, s) = (1 - |l_f - s|, 1 - |l_s - f|)$$

This is the circularity that Kagan worries is vicious.

We can re-state the problems of consistency and indeterminacy in terms of properties of the function k. A *fixed point* of k is a joint level of conformity for Frank and Stella (f, s) such that $k(f, s) = (f, s)$. Our assumptions about desert and conformity to morality are *consistent* just in case k has a fixed point. Our assumptions about desert and conformity are *determinate* just in case this function has a *unique* fixed point. Is there a unique fixed point?

Actually, we want the answer to a more general question than this one. The function g, and so the function k, depended on independently specified facts about how much Stella and Frank love each other. If we change the assumptions about how much they love each other, we get a different function (corresponding to different values of l_f and l_s in equation (1)). It is no good if some assumptions about how Frank and Stella love are consistent and determinate; what we need to do is show that any assumption about how Frank and Stella love is consistent and determinate.

The new theory I have presented so far, which allows desert and conformity to vary continuously, does not yet solve the problem of indeterminacy. Here is why. Suppose that Frank and Stella love each other maximally. So we are supposing that $l_f = l_s = 1$.

These assumptions are consistent, but they are not determinate. Just as in Kagan's simple example, there are two fixed points: the situation in which Frank and Stella both conform maximally to the moral requirement, and the situation in which they both conform minimally. (Numerically, the fixed points are $(1, 1)$ and $(0, 0)$.)

So far, then, adding continuous variation appears to be no help at all. In fact we have made things worse. There are now infinitely many fixed points: for any level of conformity between 0 and 1 it is consistent to suppose that both Frank and Stella conform to that level.

What are our options? One option is to fiddle around with the function k. But given the other assumptions I have made, this does not seem plausible. Any plausible way of defining k must leave $(1, 1)$ and $(0, 0)$ as fixed points.

Another option, one that succeeds, is to consider an idea that Kagan takes up and rejects. Suppose that the requirement to love others to the degree they deserve is only one of the moral requirements, but not the only one. Let there be others. It doesn't matter what they are. Suppose there are n of them, and that Frank and Stella both live up to each of those other requirements to some degree. (I assume that those other requirements can be met to degrees intermediate between 1 or 0, but this is not essential.) Then Frank's level of conformity to morality depends on the degree to which he has satisfied *all* of these moral requirements, not just the requirement to love Stella (and similarly for Stella's level of conformity). We have to adjust the function g to take these other facts into account.

How should we do this? If I have satisfied requirement R_1 to degree d_1, requirement R_2 to degree d_2, and so on, then what is my level of conformity, overall? Clearly the better I do on the requirements, overall, the higher my total level of conformity. And I should turn out to conform perfectly if and only if I perfectly satisfy all the requirements. One way to reflect these constraints is to have my level of conformity be equal to the average of the degrees to which I have satisfied the various moral requirements:

$$\frac{d_1 + \cdots + d_n + \text{(degree to which I love in proportion to desert)}}{n + 1}.$$

Okay, so let us assume that in addition to knowing how much Frank and Stella love each other, we know to what degree each one has satisfied the other moral requirements. Let the sum of the degrees to which Frank has satisfied them be d_f, and the sum of the degrees to which Stella has satisfied them be d_s. Then Frank's level of conformity is just

$$\frac{d_f + \text{(degree to which Frank satisfies the requirement to love Stella)}}{n+1}.$$

Let's suppose, in fact, that Frank and Stella are very conscientious when it comes to requirements other than the requirement to love. They satisfy all of those requirements to degree 1. Then $d_f = d_s = n$. The function k then becomes

$$k(f, s) = \left(\frac{n + (1 - |1 - s|)}{n+1}, \frac{n + (1 - |1 - f|)}{n+1}\right) = \left(\frac{n+s}{n+1}, \frac{n+f}{n+1}\right). \tag{2}$$

Now $(1, 1)$ is still a fixed point of this function. But $(0, 0)$ is not: $k(0, 0) = (n/n + 1, n/n + 1)$. And in fact, it can be proved that k has a *unique* fixed point. What's more, and this is the important part, it can be proved that the function got from *any* assumptions about how Frank and Stella love, and how well they meet their other moral requirements, has a unique fixed point. That a fixed point exists solves the problem of inconsistency, and that it is unique solves the problem of indeterminacy. (The proof, which is a little bit technical, is in the appendix.)

4 A Solution for Consequentialists

Is the solution I presented in the previous section available to consequentialists? The answer is "no." Briefly, this is because the solution requires there to be multiple moral requirements, only one of which involved desert. But consequentialists believe that there is only one moral requirement.

There is a solution for consequentialists, though. Rather than appealing to additional moral requirements, they can appeal to additional grounds for desert. We are interested in defending the claim that what someone deserves depends, in part, on how well they conform to the requirements of morality. This is still true even if facts about what someone deserves also depend on other factors in addition to how well they conform. Suppose, for example, that just in virtue of being alive, people deserve to be living lives that are pretty good—certainly better than lives that are just barely worth living. Something like this is all it takes to solve the problem of indeterminate desert for consequentialists.

Appealing to these additional grounds for desert is not a desperate move made just to solve this particular problem. It is something that consequentialists have wanted to say for independent reasons. (Fred Feldman [1997b], for example, uses this kind of claim to show that consequentialists can avoid Parfit's repugnant conclusion.[7])

I will sketch here what the details of this theory are; the complete presentation is in an appendix. The function g is just as it was in the previous theory. But the function h looks different. So let's work through again how the function h works. Recall what h does:

$$h(\text{degree of conformity}) = (\text{degree of love deserved}).$$

We originally assumed that h was the identity function: $h(f, s) = (f, s)$. That is the right thing to say if your degree of conformity is the only factor going in to what degree of love you deserve; but it no longer is. There is another factor—that you are alive. Let's suppose that, if the fact that someone is alive were the only factor going in to what degree of love they deserve, then they would deserve to be loved to degree b. How do we combine these two factors? We have numbers saying what degree each factor would determine were it acting on its own. To see what they do when they act together, let's think through some test cases. Let's assume that b is less than one. So the fact that someone is alive, all by itself, does not mean that they deserve complete and total love from everyone else. What effect does someone's level of conformity, then, have on their all-things-considered level of desert? To start, suppose that Frank perfectly conforms to morality. Then looking to conformity alone, Frank deserves perfect love. Looking to the fact that he is alive, he deserves imperfect love (love to degree b). A natural way to combine the impact of these factors is to split the difference, and say that he deserves imperfect love, though love of a higher degree than if we had not factored in his level of conformity at all. For another example, suppose that Frank

[7]Parfit presents the repugnant conclusion in his [Parfit 1984: chapter 17].

never does what he should, so conforms to morality to degree zero. Looking to conformity alone, Frank deserves perfect hate; looking to the fact that he is alive, he deserve imperfect love. Again a natural response is to split the difference, and say that together, he deserves imperfect love, though (this time) love of a lower degree than if we had not factored in his conformity at all. A rule that agrees with these verdicts about these examples says that when the two factors act together, we should average their inputs. The theory I am describing incorporates this rule. So in that theory h is defined like this:

$$h(f, s) = \left(\frac{f + b}{2}, \frac{s + b}{2}\right).$$

Using this h (and given facts about how Frank and Stella love) we can define, as before, a function K (capital this time, to avoid confusion) that takes levels of conformity as inputs, and returns levels of conformity. Then we can prove that this K has a unique fixed point. Since K always has a unique fixed point, there is always a single, determinate assignment of levels of conformity and levels of desert to the Frank and Stella scenario. (And Frank and Stella's levels of desert in any situation, together with the degree to which each is loved in that situation, will determine a unique value for that situation.) So consequentialists can solve the problem of indeterminate desert.

The consequentialist solution resembles the first solution I discussed. They both work by adding extra "moral factors": the first solution adds moral requirements that are independent of desert, and the consequentialist solution adds factors that contribute to desert which are independent of conformity. Then the contributions of these extra additional factors get averaged in to the contribution that conformity makes to desert, or the contribution that giving people what they deserve makes to conformity. (That we take the average plays an important role in the proofs in the appendices. It is what allows me to apply the mathematical theorem guaranteeing the existence of a unique fixed point.)

5 Remarks on the Consequentialist Solution

The consequentialist solution I described in the last section is a very limited theory. It can only be applied to simple, "toy" scenarios of the kind Kagan discusses. But my purpose has only been to show that there is a version of consequentialism that applies to those scenarios and avoids the problems of indeterminacy and inconsistency. I certainly do not think that it is the best version of consequentialism, much less the true and complete theory of moral obligation.

A more complicated, more realistic, and more plausible version of consequentialism could employ the same strategy that the simple version uses to avoid the problems of indeterminacy and inconsistency (to the extent that this problem arises for the more complicated theory). But there are two features of this strategy that one might dislike. I want to say a few words about them. (I have not been able to find a consequentialist solution that lacks these features.) The first feature I will discuss is not, I do not think, anything to worry about. But the second one is more worrying.

The first feature is the fact that the theory takes the *average* of the contribution that being alive makes to desert and the contribution that conformity makes to desert. Why, you might wonder, doesn't the theory *add* the contributions, rather than average them? The answer turns on the way the numbers that represent these contributions in the theory are to be understood. The number b represents the degree of love each person deserves in virtue of being alive, *on the assumption that desert depends on nothing else*. Similarly, f represents the degree of love that Frank deserves in virtue of his degree of conformity to morality, on the assumption that desert depends on nothing else. It would make no sense to add these numbers together to get what he deserves when both factors work.

Consequentialists who adopt this solution to the problem of indeterminate desert, then, may not be able to think about desert in the way they are used to. If a hedonistic consequentialist who adopts this solution says that, all told, I deserve 100 hedons of pleasure, we may want to ask him, how many of those hedons do I deserve in virtue of being human? And he cannot answer this question. He will instead have to say something like this: if being alive were they only relevant factor, then you would deserve 150 hedons. But you have not been living up to your obligations, and that failure to conform to morality drove the average down to 100.

But even if this is not the way we are used to thinking about desert, it is certainly an intelligible way; and we have learned in other contexts to doubt that moral factors “add together” in the way we might initially assume (see [Kagan 1988]).

I turn now to a second feature of the consequentialist solution I described. It has to do with the way the theory calculates the degree to which someone conforms to the consequentialist moral requirement. So let me first say something about how that goes.

To begin to think about this, we need to re-describe the Frank and Stella scenario. As I have been understanding that scenario, the presuppositions of the scenario are themselves inconsistent with consequentialism. I have been assuming that Frank is morally required to love Stella to the degree that she deserves. But consequentialism does not (directly) say that there is any such moral requirement. (For consequentialists, the one and only moral requirement is to perform the act that, of those open to you, produces the outcome with maximal value. (Or, if there are ties, perform one of the optimal acts.) Desert appears in their axiology, not their deontology.)

To reformulate the Frank and Stella scenario in consequentialist terms, we need an explicit statement of the way that the value of an outcome is determined. And we want this method of determination to be sensitive to facts about desert. So let us assume the following:

(the value of outcome O) =(the degree to which F gets what he deserves in O)
+ (the degree to which S gets what she deserves in O).

This is the kind of desert-sensitive axiology that some consequentialists want: other things being equal, an outcome in which Stella is wicked but prospers has lower value than one in which she is decent and prospers. (In these examples, the only way to cause someone to prosper is to love them.) Of course, a complete consequentialist axiology, one suitable for the real world, will be more complicated. But this is all we need to understand what is going on in this simple scenario.

What should we say about the degree to which Frank conforms to the consequentialist moral requirement? We want this degree to be between 0 and 1. Since (I am assuming) nothing Frank does can influence Stella's actions, his degree of conformity will depend only on the contribution to the values of outcomes made by the degree to which he gives Stella what she deserves. What the theory says is this: Frank conforms perfectly (to degree 1) to morality if he gives Stella exactly what she deserves, and conforms to degree 0 if he is as far away as possible from giving Stella what she deserves. (That is, he conforms to degree 0 if she deserves love of degree 1 and he loves her to degree 0, or vice versa.)

I do not like this feature of the theory. Again, consequentialists say that the one moral requirement is to perform the act that, of those open to you, produces the outcome with maximal value. Frank has a range of acts open to him: he can love Stella to any degree between 0 and 1. You might want a consequentialist theory of this scenario to say that Frank conforms to degree 1 if he produces the best outcome open to him (by loving Stella exactly to the degree she deserves), and that Frank conforms to degree 0 if he produces the worst outcome open to him. But that is not what the consequentialist solution I described says. For example, if Stella deserves love of degree 3/4, then Frank does the worst he can if he loves her to degree 0. (No other degree of love produces an outcome with lower value.) But the solution I describe says that in this situation he still conforms to degree 1/4. In effect, this solution says that Frank conforms to morality to some positive, non-zero degree, provided that there is *some* situation in which his love is farther away from what Stella deserves, even if no act open to him would produce that situation. (In the current example, this would be a situation in which he still loved Stella to degree 0, but she deserved loved of degree 1.)

6 Conclusion

Shelley Kagan has argued that no desert-sensitive moral theory can incorporate the idea that what people deserve depends on how well they conform to the requirements of morality. All the examples he discusses suffer from the problem of indeterminacy. I have showed that this is not so, by describing a desert-sensitive moral theory that does incorporate this idea, and then *proving* that it avoids the problem of indeterminacy. The theory I present makes use of ideas that Kagan himself considered, but puts them together in a more sophisticated way. And although the theory I propose is adapted to the simplifying assumptions that Kagan makes, the ideas in it could easily be incorporated into more realistic versions of non-consequentialism.

I also described a consequentialist solution to the problem of indeterminacy. But I admitted that this version of consequentialism has features (one in particular) that do not fit perfectly with the spirit of consequentialism.

Appendix A: Proof that k has a Unique Fixed Point

The proof is an application of the contraction mapping theorem. The contraction mapping theorem says that if X is any complete metric space, with metric d, and T is a function from X to itself satisfying

$$d(Tx, Ty) \leq q \cdot d(x, y),$$

where q is a non-negative real number less than 1, then T has a unique fixed point. (A metric space is just a set of points with a distance function d defined on it. A metric space is complete, roughly speaking, if it has "no holes.") I will show that k is a contraction mapping from M to itself.

First, I need to say how I am talking about distance in M, the set of conformity states. Instead of using the "Euclidean" distance on M, I use the "square metric":

$$d((x_1, x_2), (y_1, y_2)) = \max\{|x_1 - y_1|, |x_2 - y_2|\}. \tag{3}$$

So the distance between two points is either equal to their "vertical" separation or their "horizontal" separation, whichever is greater. M is a complete metric space using this metric.

Now to show that k is a contraction mapping. In general, k has the form

$$k(f, s) = \left(\frac{d_f + (1 - |l_f - s|)}{n + 1}, \frac{d_s + (1 - |l_s - f|)}{n + 1} \right). \tag{4}$$

It is not too hard to see that k is a contraction. Look at what happens when we apply k to two points in M, (x_1, x_2) and (y_1, y_2). Let us write $k(x_1, x_2) = (k_1(x_1, x_2), k_2(x_1, x_2))$ (and similarly for the y's); then $k_1(x_1, x_2)$ is just Frank's resulting level of conformity, giving initial the joint initial conformity state (x_1, x_2). Looking just at what happens to Frank's level of conformity, we see that

$$\begin{aligned}
|k_1(x_1, x_2) - k_1(y_1, y_2)| &= \left| \frac{d_f + (1 - |l_f - x_2|)}{n+1} - \frac{d_f + (1 - |l_f - y_2|)}{n+1} \right| \\
&= \frac{||l_f - x_2| - |l_f - y_2||}{n+1} \\
&\leq \frac{|l_f - x_2 - l_f + y_2|}{n+1} \\
&= \frac{|y_2 - x_2|}{n+1} \\
&\leq \frac{1}{n+1} \max\{|y_1 - x_1|, |y_2 - x_2|\} \\
&= \frac{1}{n+1} d((x_1, x_2), (y_1, y_2)).
\end{aligned}$$

(Step 2 uses the "backwards" triangle inequality, $||x| - |y|| \leq |x - y|$.) Similarlly, for Stella's conformity state,

$$|k_2(x_1, x_2) - k_2(y_1, y_2)| \leq \frac{1}{n+1} d((x_1, x_2), (y_1, y_2)).$$

Since $d(k(x_1, x_2), k(y_1, y_2)) = \max\{|k_1(x_1, x_2) - k_1(y_1, y_2)|, |k_2(x_1, x_2) - k_2(y_1, y_2)|\}$ this shows that

$$d(k(x_1, x_2), k(y_1, y_2)) \leq \frac{1}{n+1} d((x_1, x_2), (y_1, y_2)).$$

So k is a contraction mapping. That completes the proof that this theory of the way desert depends on moral worth avoids inconsistency and indeterminacy.

Appendix B: Proof that K has a Unique Fixed Point

The proof uses, again, the contraction mapping theorem. We first need an explicit formula for K. I said above that the function h now looks like this:

$$h(f, s) = \left(\frac{f + b}{2}, \frac{s + b}{2} \right).$$

Then g, as before, is

$$g(x_1, x_2) = (1 - |l_f - x_2|, 1 - |l_s - x_1|).$$

K, therefore, is

$$K(f, s) = \left(1 - \left|l_f - \frac{s+b}{2}\right|, 1 - \left|l_s - \frac{f+b}{2}\right|\right).$$

The proof that K is a contraction is a computation like the one for k:

$$\begin{aligned}
|K_1(x_1, x_2) - K_1(y_1, y_2)| &= \left|\left(1 - \left|l_f - \frac{x_2+b}{2}\right|\right) - \left(1 - \left|l_f - \frac{y_2+b}{2}\right|\right)\right| \\
&= \left|\left|l_f - \frac{y_2+b}{2}\right| - \left|l_f - \frac{x_2+b}{2}\right|\right| \\
&\leq \left|\frac{x_2+b}{2} - \frac{y_2+b}{2}\right| \\
&= \frac{1}{2}|x_2 - y_2|.
\end{aligned}$$

References

Arrhenius, Gustaf (2006). "Desert as Fit: An Axiomatic Analysis." In Kris McDaniel, Jason R. Raibley, Richard Feldman and Michael J. Zimmerman (eds.), *The Good, The Right, Life and Death*, chapter 1. Burlington, VT: Ashgate Publishing Company.

Feldman, Fred (1997a). "Adjusting utility for justice: a consequentialist reply to the objection from justice." In *Utilitarianism, Hedonism, and Desert: Essays in Moral Philosophy*, 151–174. New York: Cambridge University Press.

— (1997b). "Justice, Desert, and the Repugnant Conclusion." In *Utilitarianism, Hedonism, and Desert: Essays in Moral Philosophy*, 193–214. New York: Cambridge University Press.

— (1997c). "World Utilitarianism." In *Utilitarianism, Hedonism, and Desert: Essays in Moral Philosophy*, 17–35. New York: Cambridge University Press.

Field, Hartry (2008). *Saving Truth from Paradox*. New York: Oxford University Press.

Kagan, Shelly (1988). "The Additive Fallacy." *Ethics* 99: 5–31.

— (2006). “Indeterminate Desert.” In Kris McDaniel, Jason R. Raibley, Richard Feldman and Michael J. Zimmerman (eds.), *The Good, The Right, Life and Death*, chapter 3. Burlington, VT: Ashgate Publishing Company.

Parfit, Derek (1984). *Reasons and Persons*. New York: Oxford University Press.

Sider, Theodore (1993). “Asymmetry and Self-Sacrifice.” *Philosophical Studies* 70: 117–132.