How to Adjust Utility for Desert[*]

Bradford Skow

Abstract

It is better when people get what they deserve. So we need an axiology according to which the intrinsic value of a possible world is a function both of how well-off and of how deserving the people in that world are. But how should these "desert-adjusted" values of possible worlds be calculated? It is easy to come up with some qualitative ideas. But these qualitative ideas leave us with an embarrassment of riches: too many quantitative functions that implement those qualitative ideas. In this paper I will select one of these quantitative functions and defend its superiority.

Keywords: value, desert, consequentialism.

1 Motivating A Desert-Adjusted Axiology

What makes one possible world better than another (in the sense of "better" that matters morally)? Better possible worlds have higher intrinsic values. One theory, welfarism, says that the intrinsic value of a possible world is a function only of the levels of welfare enjoyed by the people in that world. (This theory uses the concept of individual welfare. Talk of someone's level of individual welfare is talk of how good that person's life is for him, of how much he has of what makes life worth living.) A simple form of welfarism works like this. In each possible world each person's life will has some numerical welfare level (relative to some choice of unit for measuring welfare). Then the (number representing the) value of a possible

---

world is just equal to the sum of the (numbers representing the) welfare levels of the people who exist in that world. This is the axiology that appears in "standard" versions of act consequentialism.

This simple form of welfarism is false. Consider two possible worlds, A and B. Each world contains just two people, Smith and Jones. The welfare levels of Smith and Jones are the same within and across worlds: in each world each of them has a welfare level of 10 units. But in world A Smith and Jones are evil. They have committed all sorts of terrible (morally terrible) acts. In world B, on the other hand, Smith and Jones are saintly. They respect morality and always act in accordance with its dictates. Welfarism says that worlds A and B are equally good. But that is wrong. In A wicked people are living good lives while in B good people are living good lives. Since it is worse for the wicked to prosper than for the good to prosper, world A is worse than world B. (This argument may be found in many places; see, for example, [Ross 1930: 138].)

Why is it worse for the wicked to prosper? Because the wicked do not *deserve* to prosper. And it is bad for them to get something they do not deserve. This is an instance of a more general truth: whenever there is a lack of fit between what people deserve and what they get the intrinsic value of the world they inhabit goes down. Any axiology that incorporates this idea is a desert-adjusted axiology.

Now defenders of desert-adjusted axiologies often make qualitative statements about how facts about desert should influence the intrinsic values of possible worlds, and look at a few examples and say which worlds in the examples a desert-adjusted axiology should say are best overall. But to do this is not yet to write down a full-blown, quantitative axiology that lets us compute the intrinsic value of any possible world. My aim is to propose and defend a quantitative desert-adjusted axiology.

(It is true that some defenders of desert-adjusted axiologies make somewhat detailed quantitative proposals about how possible worlds are to be ranked *from the perspective of desert.* (Shelly Kagan, for one, has written several papers on this topic; see, for example, [Kagan 1997] and [Kagan 2003].) But since these proposals isolate just the value of desert and ignore other things that are of value, they are limited in scope. To articulate a *complete* axiology one must say how the

2

ranking of worlds with respect to desert is to be weighed against the ranking of worlds with respect to other things that are of value. That is what I will do.)

One might worry that, in fact, there is no way to turn the somewhat vague idea that utility should be adjusted for desert into a consistent, quantitative axiology. But that is really nothing to worry about; consistent quantitative desert-adjusted axiologies exist.[1] A more serious danger is that there might be *too many* consistent theories. For suppose that there are several quantitative desert-adjusted axiologies. These theories, of course, disagree about the intrinsic values of some possible worlds. Suppose also that there is no non-arbitrary way to decide which theory is best. That, I think, would leave us in a bad spot. We see this worry expressed by, for example, Robert Nozick, in a discussion of utilitarian treatments of punishment: he wrote that a plausible version of utilitarianism will give "lesser weight to the punished party's unhappiness. One would suppose that considerations of desert...would play a role here; one would suppose this if one weren't bewildered at how to proceed, even using such considerations, in assigning the 'proper' weight to different persons' unhappiness" [Nozick 1974: 62]. But we need not be bewildered; I have found a way to proceed. It does look at first as if there are many desert-adjusted axiologies and no non-arbitrary way to decide between them. But it is possible to non-arbitrarily select one as the best.

2   Background

In order to have a complete desert-adjusted axiology one thing that we need is a theory of desert. Such a theory will tell us what people deserve, given other facts about them. Do people deserve certain things (to live relatively decent lives, for example) just in virtue of being people? Does immoral behavior lower one's

---

[1]However, the only consistent quantitative version (other than those I will present in this paper) is due to Erik Carlson, and his theory is flawed (for reasons I will get to in footnote 11). Thomas Hurka [2001] presents quantitative theories of value-from-the-perspective-of-desert, but no quantitative theory of overall value. Gustaf Arrhenius [2006], [2007] presents versions that are not quantitative. Fred Feldman [1995a] discusses several different ways to adjust utility for desert, but does not definitively endorse any one of them.

desert level? Does one deserve more now just because one has not gotten what one deserved in the past?

I am not going to try to answer these questions. I will not be presenting a theory of desert in this paper. (Discussions of some theses one might, or might not, want to include in a theory of desert may be found in [Feldman 1995b] and [McLeod 2008].) But I am going to make some assumptions about desert. Let me lay those out now.

I am going to assume that welfare is the "currency" of desert. For the purposes of assigning intrinsic values to possible worlds, the fundamental thing people deserve to have more or less of is welfare; it is a lack of fit between welfare received and welfare deserved that makes intrinsic value go down.(This is a commonly-made assumption; it is made in, for example, [Feldman 1995a], [Kagan 1997], [Kagan 2003].) I am also going to ignore facts about *when* people deserve certain welfare levels. I will assume, instead, that all there is to desert is that each person deserves to live a life with a certain over-all welfare level (where someone's over-all welfare level is the sum of their welfare levels at all times from their birth to their death).

So much for assumptions about desert; now let me say something about the structure of my theory. The theory I am going to present is "individualistic." There are two related components to the theory's individualism. First, in the theory we can speak of what each individual separately contributes to the intrinsic value of the world he inhabits. The total value of the world is just equal to the sum of the contributions of the individuals who inhabit the world. And second, what a given individual contributes does not depend on what is happening with the other individuals in that world.

The version of welfarism I briefly presented earlier is individualistic, and looking at it will help make clear what individualism is. That theory says that the contribution each person makes to the intrinsic value of a possible world is determined by (in fact equal to) just one thing: that person's welfare level (in that possible world). Clearly in this theory the value of a world is the sum of the individuals' contributions, and the contributions are independent. To figure out someone's contribution, we do not need to know (for example) how his welfare level compares to the welfare level of other people.

In the theory I will present the contribution each person makes to the intrinsic value of a possible world is determined by just *two* things: that person's welfare level and that person's desert level. We do not need to also know (for example) whether that person is closer to or farther away from getting what she deserves than other people are.[2] Once we know each person's contribution to intrinsic value, we proceed just as we did in welfarism: we sum up the contributions to get the overall value of the possible world.

A final bit of background. The theory I will present only treats positive levels of desert and positive levels of welfare. I have views about how the theory I will present can be extended to handle negative levels, but negative levels of welfare and desert pose special problems. How to handle negative welfare levels is an area of ongoing research.

What we need, then, for a quantitative desert-adjusted axiology is a function — the "contribution function" — from an individual's level of welfare and desert to that individual's contribution to overall value. Such functions can be represented by graphs, like the one in figure 1.

The horizontal axis of this graph represents someone's welfare level W, and the vertical axis represents his contribution C to the intrinsic value of the world. His desert level is indicated by the "D" on the horizontal axis: that is the level of welfare he deserves. We are trying to select, in a non-arbitrary way, a function f so that $C = f(W, D)$. This figure is a graph of the function $C = W$. That function is the function that welfarism uses for an individual's contribution to value; desert plays no role in it. It is not the function we are looking for.

The graph in figure 2 does take D into account. This function says that what someone contributes to value is just his welfare level, up until he is getting the welfare level he deserves; if his welfare level is greater than his desert level, though, he merely contributes his desert level. According to this function, giving someone more than he deserves does not make the world any better.

Now there are obviously tons of graphs one could draw on these axes, and so

---

[2]Shelly Kagan [2003] defends some principles of "comparative desert" that do look at these sorts of facts. I am unsure whether the axiology I present needs to be supplemented with comparative principles.

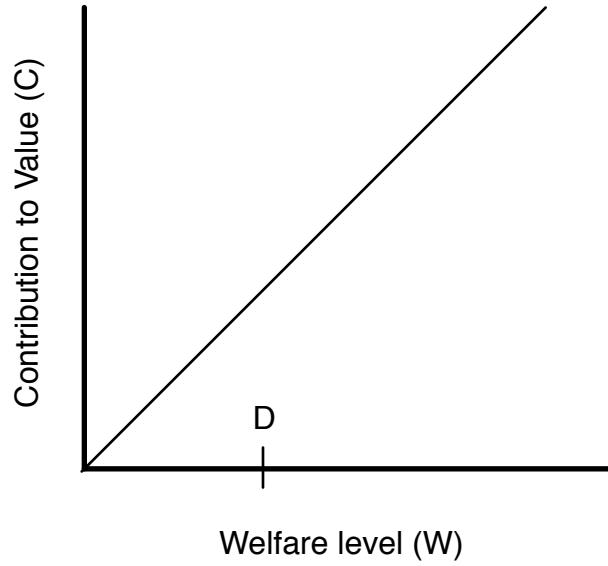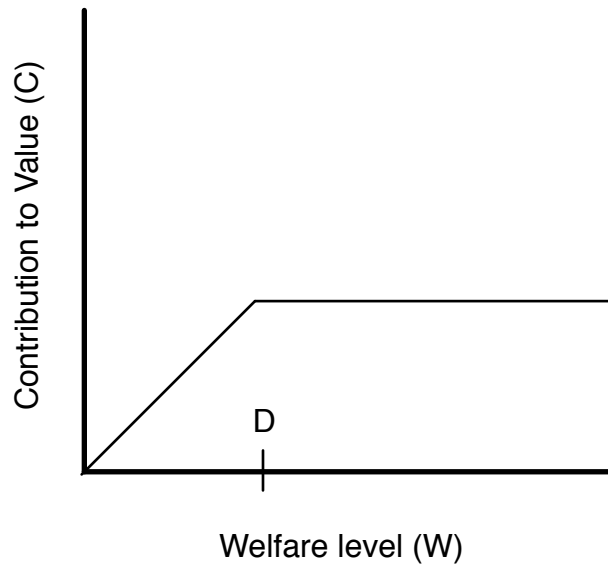Figure 1: A possible graph of Contribution as a function of Welfare and Desert

Contribution to Value (C)

D

Welfare level (W)

Figure 2: Another possible graph of Contribution as a function of Welfare and Desert

Contribution to Value (C)

D

Welfare level (W)

tons of functions that give *C* as a function of *W* and *D*. How are we to select which one to use?

3   Qualitative Restrictions

It is quite easy to come up with some qualitative (that is, non-quantitative) ideas about how *C* depends on *W* and *D*. But while ideas like this do in some sense narrow down the number of candidate contribution functions, infinitely many (uncountably many) candidates will be compatible with whatever qualitative constraints we choose. So even after we have placed qualitative restrictions on the contribution function, plenty of hard work remains. In this section I am going to state some qualitative constraints on the contribution function; the goal of the rest of the paper is to argue in favor of one of the infinitely many functions compatible with these constraints.

I will say a few things in defense of the constraints I impose. But my goal in this paper is not to mount a complete defense of them. Most other authors present and defend some qualitative constraints on the contribution function and then stop ([Feldman 1995a] is a good example of this). My interest is in defending a desert-adjusted axiology that says something more specific about the contribution function than just that it satisfies these qualitative constraints, so my main interest is in defending the part of my theory that is more specific than others.

Here are the constraints. Suppose someone is getting more than he deserves: he is living a life better than the life he deserves to live. This is a case of over-receipt. Similarly, when someone is getting less than he deserves we have a case of under-receipt. The first criterion I want the theory to satisfy is that there be an asymmetry between over-receipt and under-receipt. I want over-receipt to be less good than under-receipt is bad. To motivate this asymmetry, suppose we have two people who are at the same desert level — they deserve to be living lives of equal welfare levels — each of whom is getting what he deserves. Suppose we could redistribute the welfare, taking one unit of welfare from one of them and giving it to the other. Then one person will receive one unit less than he deserves, and the other will receive one unit more than he deserves. This situation is worse than the

one we started with. So we want our theory to say that the one unit decrease in welfare corresponds to a greater than one unit decrease in intrinsic value, while the one unit increase in welfare corresponds to a less than one unit increase in intrinsic value. Geometrically, this means that the slope of the contribution function should be smaller to the right of D than it is to the left of D.

In fact I want my theory to incorporate something stronger than an asymmetry between over-receipt and under-receipt. I want the theory to assert the "diminishing marginal utility of surplus welfare" and the "increasing marginal disutility of deficient welfare." Here is a qualitative statement of the first of these principles:

MU-1 If someone is getting more than he deserves, then each one unit increase in his level of welfare contributes less and less to intrinsic value. And the marginal utility of surplus welfare does not just diminish; it diminishes towards zero. That is, as someone gets farther and farther over his desert level, the contribution each one unit increase in his level of welfare makes gets arbitrarily close to zero.

Similarly, the second of these principles (call it MU-2) says that if someone is getting less than he deserves, then each one unit decrease in his level of welfare contributes more and more negatively to intrinsic value.[3] (Together these amount to the diminishing marginal utility of welfare generally.) MU-1 and MU-2 impose a stricter requirement on the shape of the graph of the contribution function. They require that the slope of the graph decrease as W increases (and decrease to 0 as W goes to infinity).

There is one final requirement I should make explicit (it is implicit in the statement of MU-1 and MU-2). I require the axiology to say that

(D) Increases in welfare always contribute positively to intrinsic value. It is always better, overall, to make someone better off.

Like the earlier principles, (D) corresponds to a constraint on the shape of the graph of the contribution function. It entails that the contribution function is an increasing function: the greater W is, the greater C is.

---

[3]Many philosophers have found these principles appealing. See, for example, [Feldman 1995a: 575], [Kagan 2003: 96], and [Carlson 1997: 311].

Nothing in the idea that we should adjust utility for desert forces us to accept (D). One might think that if someone is getting far more than he deserves then things are worse than they would be if he got exactly what he deserves. I, however, do not take this view. Let me say a few things about this.

The first thing to bear in mind is that I am only developing the part of the axiology that applies to positive levels of desert and welfare. It is only when these levels are positive that (MU-1), (MU-2), and (D) are true. They are not plausible for negative desert levels. (D) is certainly not plausible in that case: if someone deserves to be living a life with a negative welfare level and he enjoys a life high in welfare, it is a bad thing to make his life even better.

A second (and more important) point is that the axiology I am looking for is one that assigns "overall" or "all things considered" values to possible worlds. It does not just say how good each world is "from the perspective of desert." For there are also other values, other (morally relevant) perspectives from which worlds may be ranked in value. I am looking for an axiology that combines a world's rankings with respect to several morally relevant perspectives and gives it a single overall ranking.

In light of this fact let us look at (D). One respect in which a world may be good is that there may be a good fit in it between the welfare levels people have and the welfare levels they deserve.[4] But there is another respect in which a world may be good; a world is good in this respect when the people in that world live lives with high levels of welfare. (Or, at least, high levels of welfare are good so long as the people with those high levels have positive desert levels. Again, I am not going to address what happens when people have negative levels of desert in this paper.) Tom, let us suppose, deserves (only) a moderately good life, but he is living a *fantastically* good life. Luck and good fortune smile upon him. If we improve Tom's life the world gets worse in one respect (there is a worse fit between desert

---

[4]Kagan [2003] distinguishes the value of noncomparative desert from the value of comparative desert. He thinks that one world may be better than another with respect to noncomparative desert while being worse with respect to comparative desert. In this paper I am ignoring the value of comparative desert. So when I speak of how worlds are ranked from the perspective of desert, by "desert" I mean what Kagan means by "noncomparative desert."

and welfare) and better in another respect (there is more welfare). What happens to the overall value of the world? Does it go up or down? According to (D) the overall value goes up, even though things have gotten worse in some respect. That is how the axiology I will propose trades these values off of each other in this kind of case.

I have no knock-down argument against the alternative view: that at some point the badness of the decrease in fit outweighs the goodness of the increase in welfare. But it is important to see that the principles I have listed do incorporate the idea that there is something bad about a world in which someone gets far more than he deserves. Principle MU-1 says that the surplus welfare Tom receives contributes less to the intrinsic value of the world than does the welfare that gets him "up to" his desert level. It is this decrease in the value of his welfare that reflects the fact that increasing his welfare is bad in some respect.

I have presented some qualitative conditions on the contribution function and said something in their defense. But coming up with qualitative conditions is the easy part, for it is easy to have intuitions about which qualitative conditions are correct. The hard part remains. How could we hope to choose one of the infinitely many functions that meet these qualitative conditions? This looks like a place where intuition gives out. Some of the functions that meet the qualitative conditions are graphed in figure 3. How are we to make a non-arbitrary choice of one of these, and so a non-arbitrary choice of a desert-adjusted axiology?

4   Remarks about Incommensurability

Before going on I want to say something about incommensurability. According to the kind of desert-adjusted axiology I favor there is more than one "source" of intrinsic value. It is a good thing when someone lives a life high in welfare. And it is also a good thing when there is a close fit between someone's welfare level and their desert level. So one might wonder: are these values commensurable? For any possible world it makes sense to assign that world a value on a scale that measures amounts of welfare, and it makes sense to assign that world a value on a scale that measures degree of fit between welfare and deserved welfare. But does it

Figure 3: Functions Meeting the Qualitative Constraints
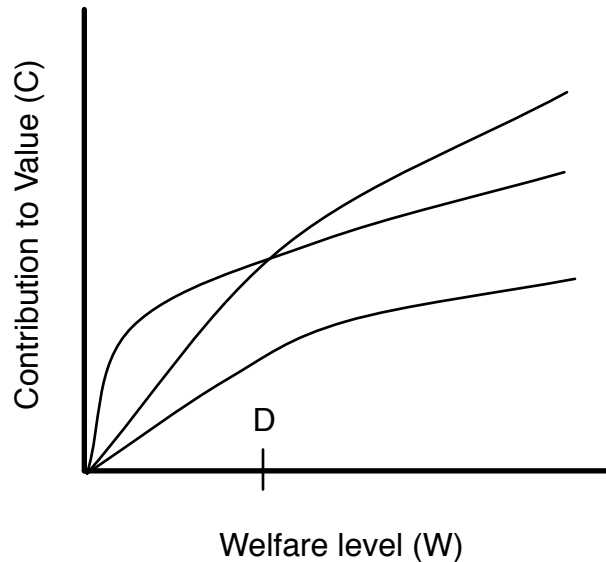
also make sense to assign it a "combined" score?[5] If it does not then my project is doomed from the start. Or again: when the fit between what someone gets and what they deserve goes down by a certain amount while their welfare level goes up by a certain amount, does the value of the possible world they inhabit go up or down? By how much? A desert-adjusted axiology attempts to answer these questions, but if these values are incommensurable these questions do not have answers.

To some extent, I think, worries about incommensurability here have the same source as the worry about arbitrariness that I articulated at the beginning of this paper. There seem to be so many ways to adjust utility for desert — to assign a world a combined score. But it can seem like there are no *non-arbitrary* ways to do it; and that suggests that there is not in fact any common scale here at all.

A defense of a desert-adjusted axiology should include a response to these

---

[5]There are several distinct things that philosophers have used "incommensurable" to express. I am using it to express what [Hsieh 2007] calls "the second conception of value incommensurability": "there is no true general overall ranking of the realization of one value [in this case, welfare] against the realization of the other value [fit between welfare and desert]."

worries about incommensurability. Part of my response is just the quantitative theory I present: seeing how the theory assigns possible worlds an overall value will (I hope) convince readers that it make sense to do so. But I can also draw attention again to the qualitative principles I articulated in the last section. They are principles about how certain changes in welfare and changes in fit together affect overall value. (The type of situation mentioned in MU-1, for example, is a type of situation in which welfare goes up but fit goes down; MU-1 says (in part) that the positive impact on overall value due to the increase in welfare is always greater than negative impact on overall value due to the decrease in fit.) The fact that those principles are not obviously misguided shows that we do not start with the intuition that welfare and fit are incommensurable.

5 Measuring Fit

Simple welfarism has a "straight line" contribution function. But the contribution function for a desert-adjusted axiology must curve. The problem I am discussing is: exactly how does it curve?

One way to get started on this problem is to find a quantitative way to measure the fit between the welfare level someone deserves and the welfare level they receive. Then we could use that measure of fit to control how the contribution function curves.

But it is not obvious how to choose a quantitative measure of fit. There are (at least) two initially appealing candidates. To see what they are, first consider a world in which one person, A, deserves 10 units of welfare and receives 20, and a second person, B, deserves 2 units of welfare and receives 4. Each of these people is getting more than he deserves. Suppose we have an additional unit of welfare to distribute. Would we make things better by giving it to A, or by giving it to B?[6] If

---

[6]It is easier to put this question in temporal terms, but officially I am ignoring facts about how welfare is distributed over time. Officially I am asking: consider two alternative scenarios, one in which A gets one more unit of welfare than he actually receives while B's welfare level stays the same, and another in which it is B who gets one more unit than he actually receives. Which scenario is better?

we accept (simple) welfarism, we do not care.[7] But if we accept a desert-adjusted axiology, we do.

The world we are considering is one in which the fit between desert and receipt is not perfect. Neither person receives exactly what he deserves. And when we distribute our extra unit of welfare we will make the fit between desert and receipt even worse: either A or B will then be even farther away from getting what he deserves. So which choice does the least violence to the fit between desert and receipt?

The problem is that there are initially plausible arguments on either side, each argument using a different measure of the fit between desert and receipt:

> **An Argument for improving A:** If we give the extra welfare to A, we bring it about that A receives 210% of what he deserves. Meanwhile, B will still receive just 200% of what he deserves. But if we give the extra welfare to B, then we bring it about that B receives 250% of what he deserves. Meanwhile, A will still receive just 200% of what he deserves. So we do the least violence to the fit between desert and receipt by giving the extra welfare to A.

> **An Argument for improving B:** A is already receiving *way* more than he deserves: 10 units more, in fact. B, on the other hand, is only receiving 2 more units than he deserves. We do the least violence to the fit between desert and receipt by giving the extra welfare to B.

The first of these arguments uses the ratio (welfare level)/(desert level) to measure the fit between desert and receipt, and the other uses the difference (welfare level) − (desert level). I do not think that intuition is more in agreement with one of these arguments than with the other. As far as this situation goes, then, intuition does not favor the "ratio measure" of the fit between desert and receipt over the "distance measure," or vice-versa. Must we just throw up our hands and say that the ideas that motivate desert-adjusted axiologies are not strong enough to help us figure out what to do here?

---

[7]More sophisticated versions of welfarism do not say this. Prioritarianism [Parfit 2000], for example, says that it is better to improve B.

Maybe there are other cases in which the use of one of these measures does fit better with intuition than the use of the other.[8] If so, maybe we could use that fact to decide what to say about this case. But even then all we will have done is to decide how best to measure the fit between desert and receipt. We would still not have decided how to use that measure to calculate the intrinsic values of possible worlds.

6 Dimensional Analysis

To begin to make progress on the problem I will use dimensional analysis. The problem is tailor-made for this technique: our problem is to find the correct equation relating C, W, and D, and dimensional analysis just is a technique (commonly used in physics) for narrowing down the number of candidate equations relating a set of quantities. Dimensional analysis will not provide a complete solution, but it will move things forward in two ways. First, dimensional analysis will give us good reason to prefer the ratio measure of fit to the distance measure. And second, dimensional analysis will reduce the complexity of the problem. It reduces the problem from that of choosing a function of two variables to that of choosing a function of just one.

What is dimensional analysis and how does one apply it? First let me explain what is meant by talk of the "dimensions" of a quantity. Any measurable quantity has an associated dimension. The distance between Boston and New York, for example, has the dimension of length (which we can abbreviate as "L"), and the

---

[8]The theory I will defend uses the ratio measure. One obvious problem with this measure is that someone who deserves 10 and gets 0 appears to be closer to getting what he deserves than someone who deserves 100 and gets 0. But the ratio measure says that both get 0% of what they deserve. This is not a problem for my theory: my theory will say that the second person contributes more to intrinsic value than the first. More difficult are cases where someone deserves 0 welfare. Then the ratio of welfare to desert is undefined. I will say something about this case at the end of the paper. (Shelly Kagan [2003] and Gustaf Arrhenius [2007] discuss these objections to the ratio measure. They also raise other objections that turn on cases in which someone deserves or receives a negative welfare level. I am not treating that kind of case in this paper.)

duration of my last birthday party has dimension of time (T). If we use meters to measure distance and seconds to measure time then the distance between Boston and New York will be some number of meters and the duration of the party will be some number of seconds. Those quantities have very simple dimensions. But some quantities have compound dimensions. The dimension of my current rate of motion is length/time (or L/T — this quantity will be measured to be some number of meters per second), and the dimension of the gravitational force the sun exerts on me is (mass × length)/time^2 (or ML/T^2, where M is the dimension of mass; I explain how one figures out the dimensions of force below).

One applies dimensional analysis to find the functional relationship between a set of quantities as follows. One starts with a quantity of interest and lists all of the other quantities on which this quantity depends. We might, for example, want to know the impact speed v of an object dropped from rest near the surface of the earth. The impact speed depends on only two other quantities: the height h from which it is dropped, and the gravitational acceleration g near the surface of the earth. We want to find the function f such that $v = f(h, g)$. To find it we use the requirement that $v = f(h, g)$ be *dimensionally consistent*. This requirement reduces the number of possibilities for f (hopefully down to something manageable).

What, then, is dimensional consistency? An equation is dimensionally consistent when the left hand side has the same dimension as the right hand side. For example, Newton's second law $F = ma$ is dimensionally consistent, but neither $F = m$ nor $F = a$ is. (The dimension of a product of quantities is just the product of their dimensions. So the dimension of ma is the product of the dimensions of m (M) and of a (L/T^2 — "meters per second squared"). This product — ML/T^2 — is also the dimension of force. So $F = ma$ is dimensionally consistent. [In fact, one can use the fact that $F = ma$ is dimensionally consistent to find the dimensions of force.] On the other hand, the right hand side of $F = m$ has dimension M and the left hand side has dimension ML/T^2, so this equation is not dimensionally consistent.)

To illustrate the use of dimensional analysis let me finish the impact speed example. In this case dimensional analysis tells us almost exactly what the function f is. The dimensions of the quantities we are interested in are these:

15

v has dimensions L/T.

h has dimensions L.

g has dimensions L/T^2 (the dimensions of acceleration).

There is only one combination of h and g that has the same dimension as v: $\sqrt{hg}$. So it follows that $v = A\sqrt{hg}$ (where A is some number ($\sqrt{2}$) that can be found from experiment, or by deriving the formula from Newton's laws).

(Where does the requirement of dimensional consistency come from? This is a large question, but I will try to give a brief answer. An equation like $v = \sqrt{2hg}$ expresses a relation among v, h, and g only indirectly, by way of expressing a relation among numbers assigned to these quantities by scales of measurement. What the equation says is that the number assigned to a body's impact speed on a certain scale of measurement is equal to the square root of twice the product of the numbers assigned to the height from which it is dropped and the gravitational acceleration. A *dimensionally consistent* equation has the property that it is true *no matter what scales of measurement we use.*[9] Of course, there is nothing to stop us from writing equations that are true only relative to one particular scale of measurement. Galileo's law of free fall relates the distance d fallen by a body dropped from rest to the time t it has been falling. The dimensionally consistent way to write Galileo's law of free fall is $d = 1/2gt^2$, where g (again) is gravitational acceleration. But we are free to write this law as $d = 16t^2$. This equation is true, so long as we measure distance in feet and time in seconds. (It is not true if we use meters to measure distance.) So the requirement of dimensional consistency is *not* the claim that all true equations are dimensionally consistent. It is, instead, the claim that any true equation that gives one quantity as a function of *all of the quantities on which that quantity depends* is dimensionally consistent. The equation $d = 16t^2$ gives d as a function of t; but d is not a function of t alone. This equation fails to be dimensionally consistent because it does explicitly include g. (Of course

_____

[9]There is actually some subtlety about what dimensional consistency means; really a dimensionally consistent equation expresses a truth relative to any scale from a certain class. For the full details, including the connection between dimensional consistency and invariance under change of scale, see [Barenblatt 1996].

the value of *g* in a particular system of measurement is used to calculate the 16.) So now: why must an equation that gives one quantity as a function of all those on which it depends be dimensionally consistent? Suppose we have an equation that is not dimensionally consistent. It is true only relative to a special scale of measurement. But Nature does not select any scale of measurement as special. So if the equation is true only relative to one special scale of measurement, what makes that scale special must be that it implicitly represents some quantity on which the quantity of interest depends, a quantity that has not been represented explicitly in the problem (as $d = 16t^2$ does not explicitly represent *g*). That is, the equation does not give the quantity of interest as a function of all the quantities on which it depends. And that is what was to be shown.)

Now we can get back to axiology. The requirement of dimensional consistency puts constraints on the form of the contribution function. But before we can use dimensional analysis we need to know what the dimensions of *C*, *W*, and *D* are. This is not at all obvious. These quantities certainly do not have any of the dimensions we are familiar with from physics — length, time, mass, charge, and so on. They must have dimensions that it is the special province of value theory to investigate.

If two quantities have the same dimension, then they can be measured on a common scale. All quantities with dimension length, for example, can be measured in meters. Welfare and deserved welfare can certainly be measured on a common scale, so they have the same dimension. Let us call this the dimension of welfare (which I abbreviate as "*E*" — I am already using "*W*" to denote the value of an individual's welfare).

It seems false, though, that the intrinsic value of a possible world can be measured on the same scale as the welfare level of a given individual. So the dimension of this quantity is another new fundamental dimension, which I shall call the dimension of value (*V*).

Now dimensional analysis tells us that there is a dimensionally consistent equation giving *C* as a function of all the quantities on which it depends. The axiology I am interested in certainly says that *C* depends on *W* and *D*. Does it depend on anything else? The answer, interestingly, turns out to be "yes." That is

17

because there is *no* equation relating C, W, and D that is dimensionally consistent. No combination of W and D has the same dimension as C. The dimension of any such combination is some power of the dimension of welfare: the dimension of DW is E^2, the dimension of W^5/D^2 is E^3, and so on. So even the simple equation

$$C = W \qquad\qquad (1)$$

is not dimensionally consistent: the left hand side has dimension V while the right hand side has dimension E.

The situation here is analogous to one we find in physics: the internal energy of an ideal gas is proportional to its temperature. But temperature and energy have different dimensions (in the SI system of units), so there is no dimensionally consistent equation relating just them. That is why Boltzmann's constant is there: Boltzmann's constant k_b has dimensions (Energy/Temperature), so internal energy U and k_bT have the same dimension. Thus the energy equation for an ideal gas, $U = 3/2Nk_bT$ (N the number of molecules), is dimensionally consistent.

If our problem is to have any solution, then, there must be an "evaluative constant" that plays a role in the axiology analogous to the role played by Boltzmann's constant in statistical mechanics. This evaluative constant — which I will call "k" — will have dimension V/E.[10]

Are there any *other* quantities, besides W, D, and k, on which C depends? I do not think so. What could they be? Additional fundamental evaluative constants? I cannot see how we could ever have reason to believe that there are any of those.[11]

---

[10]Earlier I wrote $C = W$ and said it was the contribution equation for simple welfarism. I am now saying that the equation should be $C = kW$. I wrote the dimensionally inconsistent equation earlier because I did not want to complicate that discussion with the demands of dimensional consistency.

[11]Erik Carlson's [1997] equations for a desert-adjusted axiology are not dimensionally consistent. One of Carlson's equations is $C = D + \sqrt{W - D}$ (for situations where $W \geq D$ — Carlson leaves it open whether to take the square root or, say, the cube root). D has dimension E while $\sqrt{W - D}$ has dimension $\sqrt{E}$, so the right hand side of this equation does not even have a well-defined dimension. So Carlson's theory presupposes the existence of some new fundamental evaluative constant (call it b) with dimension $\sqrt{E}$; then the dimensionally consistent expression of his theory

Can dimensional analysis select for us a single, unique equation giving C as a function of D, W and k? Dimensional analysis does let us rule out some candidates. But (unfortunately) it does not work the magic it worked in the impact speed example. It does not select a unique function. There are many dimensionally consistent equations relating C, W, D, and k. Let us look at some of them.

One equation is

$$C = k\sqrt{DW}. \tag{2}$$

C has dimension V; the product DW has dimension E^2, so its square root has dimension E, so $k\sqrt{DW}$ has dimension V.

Another equation is

$$C = kD\left(\frac{W}{D}\right)^{1/3}. \tag{3}$$

In fact, for any number q it is dimensionally consistent to say that

$$C = kD\left(\frac{W}{D}\right)^{q}. \tag{4}$$

More generally still, for any function (from numbers to numbers) g it is dimensionally consistent to say that

$$C = kD\,g\left(\frac{W}{D}\right). \tag{5}$$

Equation (5) is important: all of the earlier equations are instances of (5), and it can be proved that *every* dimensionally consistent equation relating C, W, D and k has the form (5).[12] So even though dimensional analysis does not select a unique

---

is $C = k(D + b\sqrt{W - D})$. Even though Carlson's equation satisfies the qualitative constraints from section 3, I do not think it is true, because I do not think there is any such quantity as b.

[12] The Buckingham Pi Theorem tells us that each equation relating C, W, D and k is equivalent to some equation relating two independent dimensionless groups constructed from C, W, D and k. C/kD and W/D are two independent dimensionless groups, so each dimensionally consistent equation is equivalent to some equation of the form $C/kD = g(W/D)$, where g is an arbitrary function. (See chapter 1 of [Barenblatt 1996] for an explanation and proof of the Buckingham Pi Theorem.)

Of course, C/kW and D/W are another two independent dimensionless groups, so each dimensionally consistent equation is also equivalent to some equation of

candidate equation, we have made progress: we know the general form the equation must take. And we have reduced the complexity of the problem: instead of needing to select a function of two variables (D and W) that meets the qualitative constraints, we only need to select a function of one variable (the function g).

Since the equation we seek must have the form (5), we need to choose what the function g shall be. Which should we choose? Before answering that question, I want to pause to point out something else that dimensional analysis has taught us about our problem. Notice that the quantity W/D appears in (5). This is just the ratio of welfare level to desert level: it is the proportion of what someone deserves that they receive. Importantly, this quantity is *dimensionless*. It does not change its value when we change our unit for measuring welfare. (Using one unit maybe I deserve 10 and receive 5; using another I deserve 100 and receive 50. Either way $W/D = 1/2$ — I am getting half of what I deserve.) So we can apply any function (from numbers to numbers) to W/D we like; when we then multiply the result by kD, we have a quantity with the dimension V.

Dimensional considerations, then, suggest that we use the quantity W/D to measure the fit between someone's welfare level and their desert level. Dimensional considerations do not quite force us to use W/D. (If g is the function $g(x) = 2 - x$, then $C = kD(2 - W/D) = k(D + (D - W))$, and this equation can be interpreted as using $D - W$ (distance) to measure fit.) But no dimensionally consistent equation that meets the qualitative constraints in section 3 can use distance to measure fit. And, in fact, if we combine the restrictions given by dimensional analysis with those restrictions on the shape of the graph of our function we can rule out additional candidates.

I said that the slope of the graph of the contribution function must be always positive but must decrease to 0. Now the slope of the graph of the contribution

the form $C/kW = g_1(D/W)$, where g_1 is an arbitrary function. I choose the groups I did because W/D has a natural interpretation as the fit between welfare and desert.

I have only looked at equations in which C is equal to some product of powers of D and W. So perhaps it is worth mentioning here that dimensional analysis has not ruled out simple equations like $C = k(D + W)$; this equation is an instance of (5) with $g(x) = 1 + x$. This equation, however, is ruled out by the qualitative constraints (it does not satisfy MU-1).

function f is just the derivative $\partial f / \partial W$, and in light of equation (5) the qualitative constraints on the derivative of g are the same as those on $\partial f / \partial W$. (I am assuming, then, that C is a *differentiable* function of D, W, and k.) So g' must be positive and decrease to 0. This eliminates all the instances of equation (4) where $q \geq 1$.

Of course, equations (2) and (3) (among others) satisfy these conditions. Figure 4 contains graphs of these two functions (for $D = 5$). They both look pretty good. How are we to choose between them?
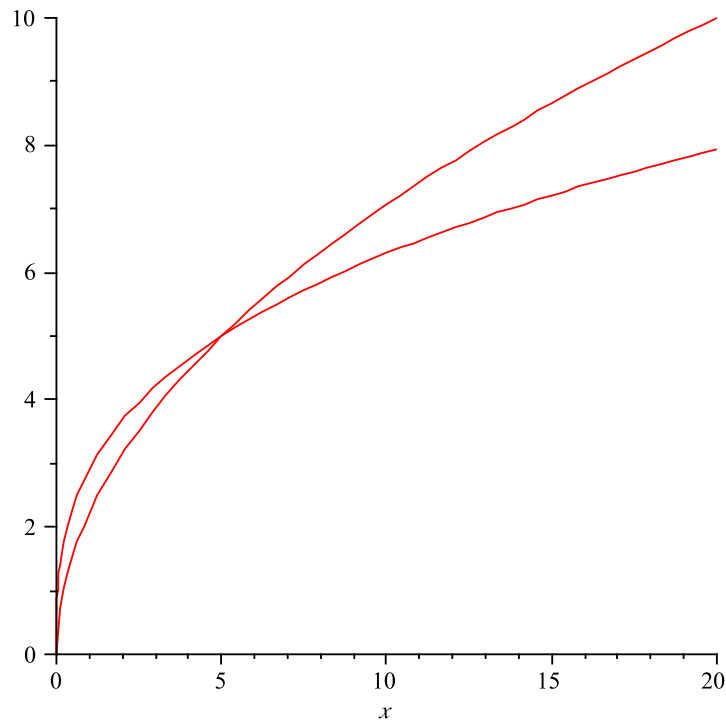

Figure 4: Graphs of equations 2 and 3

I claim that of the functions that remain, one of them is simplest. That it is simplest is a non-arbitrary fact about it; so by selecting it to appear in our desert-adjusted axiology we have made a non-arbitrary decision. (I am not going to rest my case for this function on its simplicity and non-arbitrariness; more substantive defense will come in the next section.)

Now the simplest function defined on the positive numbers that is always positive but decreases to zero is $1/x$. So a first thought is to set $g'(x) = 1/x$. This

makes $g(x) = \log(x)$, and therefore $C = kD\log(\frac{W}{D})$. But this in fact does not work, because it makes C negative if W/D is less than 1. The simplest choice that avoids this problem is to set $g'(x) = 1/(1 + x)$. Then $g(x) = \log(1 + x)$ and

$$C = kD\log\left(1 + \frac{W}{D}\right). \tag{6}$$

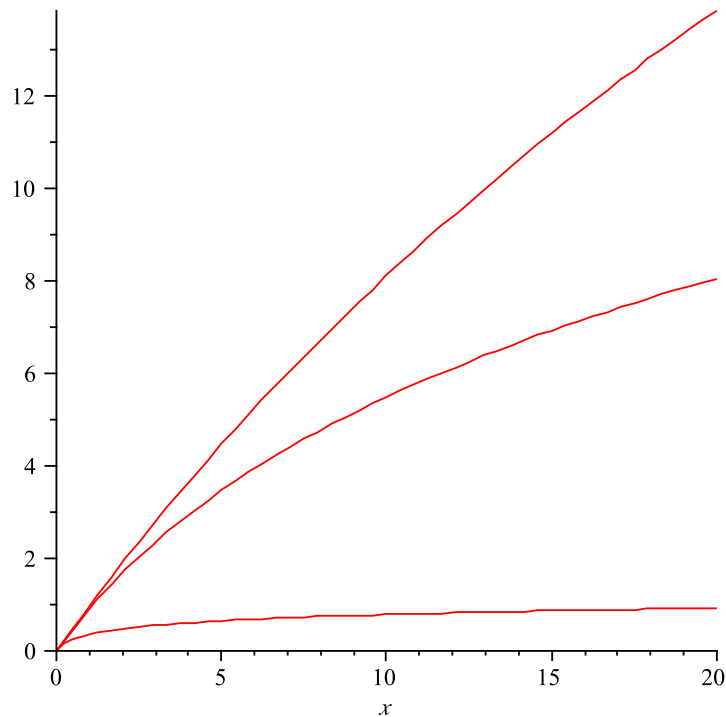Figure 5 shows three graphs of equation (6), for $D = 20$, $D = 5$, and $D = 1/5$.

Figure 5: Graphs of equation (6)

Here is something else that suggests that (6) is the right equation to use. Suppose that I deserve 10 units of welfare, and I get 10 units. We want each increase in my welfare level to be worth less in terms of intrinsic value. But how much less? Suppose I receive some additional welfare, so that my welfare level increases from 10 to 11. How much should that extra welfare contribute to the intrinsic value of the situation? That additional welfare makes it so that I am getting 11/10 of what I deserve. The simplest thing to do is to take the reciprocal of this proportion (which is

22

a number that is less than one) and multiply. That is, we say that the "discount rate" for this welfare is $(10/11)k$, and so its contribution to intrinsic value is $(10/11)k \times (1$ unit$) = (10/11)k$ units. If I receive yet another unit of welfare, increasing my welfare level to 12, then that second unit will contribute only $(10/12)k$ units to intrinsic value. And so on.

In general, suppose that I deserve D units of welfare and I receive W_s units of surplus welfare. Then the simplest proposal is that the contribution to the intrinsic value of the situation that the surplus welfare makes is

$$\left(\frac{D}{D+1} + \frac{D}{D+2} + \cdots + \frac{D}{D+W_s} \right) k \tag{7}$$

units.

I have been talking about how much the first unit of surplus welfare is worth. But how much is the first half-unit of surplus welfare worth? Surely not 1/2 the value of the first unit of surplus welfare. The first half-unit of surplus welfare should contribute more to intrinsic value than the second half-unit. (That is because surplus welfare has diminishing marginal utility.) And the first quarter-unit of surplus welfare should contribute more than the second quarter, and so on.

To achieve this result, we must use calculus and integrate rather than add. Then the equation that this line of thought suggests for the contribution that surplus welfare (C_s) makes to intrinsic value is:

$$C_s = \int_0^{W_s} \frac{kD}{D+x} \, dx. \tag{8}$$

Doing the integral we get

$$C_s = kD \log \left(1 + \frac{W_s}{D} \right). \tag{9}$$

This equation for C_s resembles equation (6) for C. Of course the equations cannot both be correct: If (6) is correct then the contribution made by the surplus welfare is

$$C_s = kD \log \left(1 + \frac{W_s}{2D} \right). \tag{10}$$

23

So the procedure that leads to (9) is not in fact correct on my theory. There is no direct argument from that procedure to the claim that (6) is correct. But I still think that the similarities between (9) and (6) are encouraging. Reflecting on the simple procedure reinforces the idea that the logarithm function is the right one to choose for g.

(Shouldn't we look for the theory according to which the procedure that leads to (9) *is* correct? The answer is that *there is no such theory*. If we try to extend the procedure to calculate the contribution of someone's initial welfare (the welfare that gets the person up to his desert level), we arrive at the absurd result that his contribution to value is infinite, no matter what his welfare level. So the fact that my theory fails to validate this procedure is not a mark against it.)

For the final version of my theory I want to modify equation (6) just a little bit. Equation (6) says that when someone's welfare level is equal to their desert level their contribution to value is $kD \log 2$, slightly less than kD. But that is a rather arbitrary thing to say. It is better to "normalize" the equation so that in this case the person's contribution to value is kD. To do this I insert a factor of $e/2$. The final equation, then, is

$$C = kD \log \left(\frac{e}{2} \left[1 + \frac{W}{D} \right] \right). \tag{11}$$

This normalization changes the graph of C a little bit; figure 6 shows what it looks like for $D = 5$ and $D = 20$.

In the next section I will give a more substantive defense of my theory. But first I want to explain why I have said so much about how I arrived at the theory (why, in particular, I have said so much about dimensional analysis). I could, of course, have just written equation (11) on the first page of this paper and said "here is my theory." I explained where it came from in order to put us in a position to see what the possible alternatives to (11) are, and to see that in selecting (11) from those alternatives I am not making a senseless and arbitrary choice.

7   Remarks on and further defense of my desert-adjusted axiology

I have used simplicity and naturalness as criteria for selecting (11) as the contribution function for my desert-adjusted axiology. The simplicity of this choice does
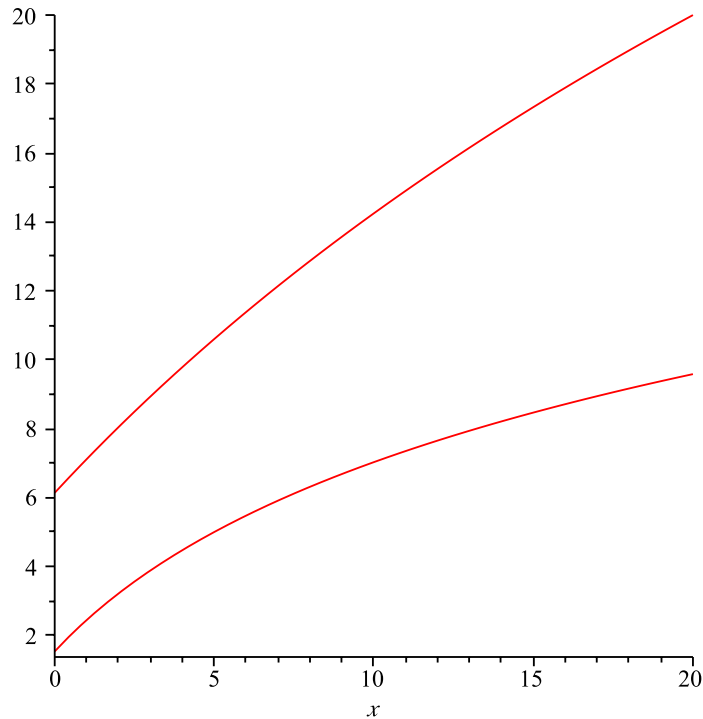
Figure 6: Graphs of equation (11)

shows that the choice is not made arbitrarily — we are not, in Nozick's words, "bewildered at how to proceed." But the simplicity of (11) is not a particularly strong argument in favor of this axiology. So in this section I will explore what this axiology says in more detail and present additional grounds for believing it.

We now know what the mathematical formalism of my desert-adjusted axiology looks like. But this desert-adjusted axiology is a part of a theory of value, not an exercise in pure mathematics. What does it feel like to use the theory?

Let us look again at the puzzling case from the beginning of the paper. Person A deserves 10 units and receives 20, and person B deserves 2 and receives 4. We have an additional unit of welfare to distribute. Who should get it? I looked at one argument that it should go to A and another argument that it should go to B. My theory says that it is better for person A to get the additional welfare. A simple computation shows that the scenario in which person A gets the additional welfare has value $10 \log(e/2[1 + 21/10]) + 2 \log(3e/2)$, which is approximately 17.19 (I

am using units where $k = 1$). But the scenario in which B gets the additional welfare has value $10 \log(3e/2) + 2 \log(e/2[1 + 5/2])$, which is approximately 17.17. Of course, once we have decided to use the ratio W/D to measure the fit between desert and receipt we can see that the argument that appealed to that measure is the good argument, and so that it is better that A get the additional welfare. What may be surprising is that the scenario in which A gets the additional welfare is *only just barely a tiny bit better* (relative to the size of our unit for measuring welfare). That is not a judgment we were in any position to make before.

I turn now to some additional grounds for believing my theory. What I will do next is look at some consequences of the theory. The consequences I will look are claims about how utility should be adjusted for desert that did not guide the construction of the theory. So we can use these consequences to test the theory: if a consequence is intuitively correct, that is evidence for the theory (and of course if a consequence is intuitively incorrect, that is evidence against it).

I set out to construct a theory in which someone's contribution to intrinsic value increases as his welfare level increases and his desert level is held fixed. (This is principle (D).) But what happens when someone's *desert level* increases while his welfare level is held fixed? Are things getting better or worse overall?

The theory says that things are getting better. That certainly sounds like the right thing to say when someone starts out getting more than he deserves. Then increasing his desert puts him closer to getting what he deserves. But the theory also says it is good to increase someone's desert level when he is getting less than he deserves. That sounds counterintuitive, so let me say something in its defense.

I will start by pursuing in more detail the intuition that my theory is wrong here. Early in the paper I discussed two things that have value: a close fit between desert and receipt is good, and welfare is good. And I said something about what the theory says about overall value in some cases where these two values conflict. Now keep these two values in mind and think about what we should say about the following case. We have someone who deserves an excellent life. But his life is only moderately good; it is not as good as the life he deserves. If his desert level goes up then the fit between what he deserves and what he receives gets worse: his welfare level was already below the level he deserved, and now it is even farther

26

below. Meanwhile his welfare level has not changed. Doesn't this mean that things must be worse *overall* — the opposite of what the theory says? No, it does not. That is because these two values are not the only two values that my axiology takes into account. It is a good thing when people lead lives high in welfare; it is a good thing when people get what they deserve; and also *it is a good thing when welfare goes to highly deserving people*. While increasing someone's desert level does not make things better with respect to either of the first two values, it does make things better with respect to this third value.

It is not hard to motivate this third value. Suppose Smith's welfare level is 10, and so is Jones's. Smith deserves 40 and Jones deserves 80, so neither is getting as much as he deserves. Doesn't it feel more urgent to improve Jones's situation? But improving Jones is *not* the best way to improve the fit between what people deserve and what they get. (Giving 10 more units to Smith gives Smith 1/2 of what he deserves and Jones 1/8. Giving 10 more units to Jones gives Jones 1/4 of what he deserves and Smith 1/4. So on the most natural way to combine these "fit" scores, the total fit gets better when the welfare goes to Smith.)

So increasing someone's desert level when he does not have as much as he deserves is good in one way and bad in another. A theory of overall value needs a way to resolve this conflict. My theory says that increasing someone's desert level while his welfare level is held fixed always makes things better overall — the good aspect of the change always outweighs the bad. Now I do not have strong intuitions about how this conflict of values should be resolved. For those who share my (lack of) intuitions the fact that my theory has this consequence is neutral: it counts neither for nor against the theory. Instead it is "spoils to the victor": if the theory gives intuitively correct verdicts about other kinds of cases, we should accept what it says about these cases as well.

Let us look at another couple of consequences of my theory. Consider first the following kind of situation. Suppose there is just one person and that he is getting more than he deserves. (Maybe his welfare level is 4 units while his desert level is merely 2 units.) Suppose that we can either increase his desert level by 1 unit or increase his welfare level by 1 unit. Which option is better? The theory says that either option will increase intrinsic value. But will one of the options increase

intrinsic value more than the other?

First consult your intuitions. It seems to me that if someone is getting more than he deserves, it is better to increase his desert level than to increase his welfare level. After all, think about this case from the first person point of view. If I were getting more than I deserved, I would feel unworthy of the life I was leading. I would devote my energy to being more deserving, not to living an even better, even more undeserved life.[13] My theory agrees.[14]

Now let us consider the opposite kind of situation. Suppose there is just one person and that he is getting less than he deserves. Is it better to increase his welfare level or his desert level (by the same amount)? Here intuition says that it is better to increase his welfare level. For increasing his welfare level is "win-win": we improve his life and we also improve the fit between what he gets and what he deserves. Increasing his desert level, on the other hand, is "win-lose." My theory agrees: although increasing his desert level does increase intrinsic value, increasing his welfare level increases intrinsic value more.

There is one more consequence of my theory that I want to discuss. I began my analysis by assuming that everyone deserves and receives some positive amount of welfare. But it turns out that my theory also tells us what we should say about a situation in which someone deserves zero welfare. This might come as a surprise because the theory uses the ratio W/D to measure the fit between receipt and desert, and W/D is undefined when $D = 0$. The theory appears to break down at this point. But it does not. Instead the theory says that people who deserve a welfare level of

_____

[13]My attention was drawn to this intuition by the discussion at UNC-Chapel Hill. Geoff Sayre-McCord was especially helpful.

[14]We want to compare $\partial C/\partial D$ to $\partial C/\partial W$. For simplicity assume that the constant k is 1. Then

$$\frac{\partial C}{\partial D} = \log\left(\frac{e}{2}\left[1 + \frac{W}{D}\right]\right) - \frac{W}{W + D}, \text{ and}$$
$$\frac{\partial C}{\partial W} = \frac{D}{W + D}.$$

So if $W > D$ then $\partial C/\partial D > 1 - \frac{W}{W+D} = \frac{D}{W+D} = \partial C/\partial W$. The case considered in the next paragraph is one in which $W < D$; in that case a similar argument shows that $\partial C/\partial D < \partial C/\partial W$.

0 are "utility black holes." No matter how high their welfare level, they contribute 0 to intrinsic value. Let me explain why equation (11) has this consequence.

For reasons I just gave, equation (11) makes no sense when $D = 0$. But the equation does approach a limit as D goes to 0; in the limit the equation is $C = 0$. (This is suggested by the graphs in figure 6: the graphs for lower values of D are much flatter than those for higher values.) So if we want the case where $D = 0$ to be continuous with what the theory says when $D > 0$, we must say that no matter what the welfare level of someone with 0 desert, that person contributes 0 to the intrinsic value of the situation he inhabits.

I can say more about what this "continuity requirement" comes to. Suppose we do not want to say that people who deserve a welfare level of 0 are utility black holes. We will say, instead, that someone who deserves a welfare level of 0 and enjoys a positive welfare level of W nevertheless contributes some positive amount, say N units, to intrinsic value. Now equation (11) says that the following can happen. There is a person with some positive level of desert; that person also receives W units of welfare; but that person contributes fewer than N units to intrinsic value. But that is absurd: surely if two people are at the same welfare level, then the one who deserves the lower welfare level must contribute less to intrinsic value than the person who deserves the higher welfare level! (This consequence is not particular to equation (11). *Every* dimensionally consistent equation that meets the constraints in section 3 must also say that people who deserve 0 welfare are utility black holes.[15])

What are we to make of the fact that my theory has this consequence? I have to admit that I am not entirely sure.[16] This consequence may appear to conflict

---

[15]We know that all dimensionally consistent equations have the form $C = kD\,g(W/D)$, where g is an arbitrary function (which I assume is differentiable). So we want to find the $D \to 0$ limit of the right-hand sides of equations with this form. Now $g(x)$ is increasing; if it approaches some finite limit as $x \to \infty$, then it is immediate that $kD\,g(W/D) \to 0$ as $D \to 0$. If, on the other hand, $g(x) \to \infty$ as $x \to \infty$, we have an indeterminate form. Then l'Hopital's rule tells us that the $D \to 0$ limit of our target equation is equal to the $D \to 0$ limit of $kW\,g'(W/D)$. But from section 3 we know that $g'(x) \to 0$ as $x \to \infty$.

[16]For what it is worth, on pages 576-577 of "Adjusting Utility for Justice" Feldman mentions the view that people who deserve *negative* levels of welfare and receive a higher level of welfare than they deserve contribute 0 units to intrinsic

with the principle (D) I gave as a constraint on desert-adjusted axiologies earlier in this paper. That principle said that increasing someone's welfare level always increases his contribution to intrinsic value. But that principle was supposed to be plausible only for people who deserved some positive amount of welfare. It is certainly not plausible for people who deserve negative welfare levels (a kind of situation my theory does not address). Still, that this consequence is consistent with the qualitative principles I listed does not make it plausible. I think, in fact, that we are not yet in a position to evaluate this feature of my theory. Whether it is acceptable to say that people with desert level zero are utility black holes depends (at least in part) on what someone has to do, how evil someone has to be, to earn a desert level of zero. But I have not said anything in this paper about what determines what desert levels people have. What we really need is a complete version of my theory, one that addresses this topic and also says what happens when someone's desert level or welfare level is negative. I hope to develop a more complete version of the theory in the future.

8 Conclusion

The simple form of welfarism I began with has a precise, quantitative axiology. Assuming we can assign numbers to represent the welfare levels of individuals in possible worlds, it allows us to compute the intrinsic value of any possible world. This axiology is false but simple; more plausible axiologies are more complicated and (so) are rarely put in the same quantitative form. Perhaps this is because it seemed too hard to put them in this form, or because there seemed to be too many arbitrary choices to be made to put them in this form. I have tried to show that this is not correct: the qualitative features we want a desert-adjusted axiology to have allow us to select one quantitative axiology as the simplest and most natural one to use. And the theory so selected has several other intuitively desirable features.[17]

_____

value. This view does not seem to strike him as implausible. However, Feldman explicitly rejects the idea that people who deserve welfare level 0 are utility black holes.

[17]Thanks to Dave Baker, David Manley, Sarah Moss, and Erik Swanson for helpful feedback, and for proposing some interesting alternatives to my view that I have

References

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

— (2007). "Meritarian Axiologies and Distributive Justice." In T. Ronnow-Rasmussen, B. Petersson, J. Josefsson and D. Egonsson (eds.), *Hommage a Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz.* `http://www.fil.lu.se/hommageawlodek/site/papper/ArrheniusGustaf.pdf`.

Barenblatt, G. I. (1996). *Scaling, self-similarity, and intermediate asymptotics*. New York: Cambridge University Press.

Carlson, Erik (1997). "Consequentialism, Distribution and Desert." *Utilitas* 9: 207–216.

Feldman, Fred (1995a). "Adjusting Utility for Justice: A Consequentialist Reply to the Objection from Justice." *Philosophy and Phenomenological Research* 55: 567–585.

— (1995b). "Desert: Reconsideration of Some Received Wisdom." *Mind* 104: 63–77.

Hsieh, Nien-he (2007). "Incommensurable Values." In *The Stanford Encyclopedia of Philosophy*. Available online at `http://plato.stanford.edu/entries/value-incommensurable/`.

Hurka, Thomas (2001). "The Common Structure of Virtue and Desert." *Ethics* 112: 6–31.

Kagan, Shelly (1997). "Equality and Desert." In Louis Pojman Owen McLeod (ed.), *What Do We Deserve?*, 298–314. New York: Oxford University Press.

— (2003). "Comparative Desert." In Serena Olsaretti (ed.), *Desert and Justice*, 93–122. New York: Oxford University Press.

McLeod, Owen (2008). "Desert." *The Stanford Encyclopedia of Philosophy* `http://plato.stanford.edu/entries/desert/`.

Nozick, Robert (1974). *Anarchy, State, and Utopia*. Basic Books.

Parfit, Derek (2000). "Equality or Priority?" In Matthew Clayton and Andrew Williams (eds.), *The Idea of Equality*, 81–125. New York: Palgrave.

Ross, Sir William David (1930). *The Right and the Good.* Oxford: Oxford University Press.