M.A. Roberts¹

What Is the Right Way to Make a Wrong a Right?

It seems clear that the most challenging versions of the nonidentity problem involve, at least implicitly, claims about probability. Once we realize that, we are tempted to appeal to the concept of *expected utility* for purposes of understanding the problem and analyzing the underlying cases. But there are reasons to think that that approach is ultimately unsatisfactory. Thus the question remains open just how probabilities are to be brought to bear in connection with nonidentity. This paper explores some of our options and some of the challenges those options will face.

¹ Philosophy, Religion, and Classical Studies Department, The College of New Jersey, robertsm@tcnj.edu.

1. Trading Off the Better Outcome Against the Better Chance

It is plausible to think that not all choices that end badly—end, for example, in outcomes in which a particular person has less wellbeing than that same person could have had at no cost to anyone else—are wrong. Consider the case of *Fertility* (Graph 1).

Graph 1. Fertility

		0.0001	0.0099	0.99
c1 potential mom doesn't undergo fertility treatment	+10 +8 +0	<i>o1</i> p	<i>o2</i>	<i>o3</i>
c2 potential mom undergoes fertility treatment	+10 +8 +0	<i>o3</i>	<i>o4</i> p	<i>o6</i>
		<i>p</i>		<i>p</i>

In this graph and the following, the acts, or choices, c1, c2 and so on exhaust the agent's (agents') available choices and the possible worlds, futures or outcomes, o1, o2 and so on exhaust the outcomes accessible under those choices.² The numbers 0-1 in the second row represent the probabilities, based on information available to the agent(s) just prior to choice, that a given outcome will unfold under a given choice. The numbers in the second column represent (overall, lifetime; raw, unadjusted) wellbeing levels. (It's left open whether wellbeing consists of happiness, capability or something else entirely.) Positive numbers represent a life worth living and negative numbers a life less than worth living. *p*, *q* and so on are constants reserved for persons, with the term person being understood to include many non-human animals and not all humans. A personal constant in bold means that the person does or will exist, and in italics that the person never exists, in the particular outcome. It is assumed that a person's wellbeing level is zero in any outcome in which that person never exists (no wellbeing burdens; no wellbeing benefits) and stipulated that no one other than the persons identified in the graph is affected (in terms of existence, wellbeing or any other potentially morally relevant respect) however the choice under scrutiny is made. A gray background indicates the choice that is in fact made in the particular case and the outcome that in fact unfolds under that choice.

² I use the term "accessibly could have had," rather than simply "could have had," in recognition of the fact that not all logically possible outcomes, however wonderful, are outcomes morality obligates agents to bring about. The fact that it's logically possible that one choice leads to a better outcome than another choice doesn't, we think, imply that the other choice is wrong in the case where the better outcome is barred by, e.g., the laws of physics.

In Fertility, a woman, the agent, must choose between undergoing an effective and relatively safe treatment for her infertility or declining that treatment. The woman in fact chooses treatment—she chooses c2 over c1—and outcome o4 in fact unfolds and the child p in fact comes into existence. But just *prior* to choice the not-yet-existing p faces an across-the-board profound existential risk: c2 makes p's existence *more* likely than c1 does, but neither c1 nor c2 *assures* p's future existence or even makes that future existence at all *likely*. All that is likely and is indeed *assured* is that, if p does come into existence—if, that is, o1 or o4 obtains—then p will have an overall, lifetime wellbeing level of +10 under c1 or +8 under c2. c2 thus *caps* p's wellbeing level at +8; under c1 and c1 alone is there *any* chance that p's wellbeing will be *maximized* at +10.

But we don't think c2 is *morally wrong*. Rather, we think c2 is perfectly *permissible*. It's true that *but for* the probabilities—the probability being *relatively* high (though not *high*) that p will exist under c2, that is, that o4 will unfold given c2, and the probability being relatively low (indeed, *low*) that p will exist under c1, that is, that o1 will unfold given c1—we *would* want to say that c2 is wrong and that c1 is obligatory. The probabilities being as they are, however, we consider c2 is permissible.

In this case and many others, the probabilities have served to convert an otherwise wrong choice into a permissible choice; they have converted a *wrong into a right*. We can surmise that the applicable moral principles generating the result of permissibility function to *trade off* a better *outcome* for a given person against a better *chance* of that person ever existing at all. Because p has a better *chance* of existing under c2, we say that c2 is permissible even though p's *wellbeing* level in o4 is less than what p *accessibly could have had* (at no cost to anyone else) in o1 under c1.³

To show that the distribution between wellbeing and chance we see in Fertility can arise in real life, we need just add some details to the case. Suppose that the woman is suffering from an elevated prolactin level, a condition that reduces the woman's chances of ovulation and thus her chances of conceiving a child. To lower her prolactin level, she has the option of taking the drug bromocriptine. But drugs often come with unwanted side effects, and let's now just stipulate (the actual medical reports on the side effects of bromocriptine being mixed) that bromocriptine leaves any child it enables a woman to conceive with adverse skin and neurological conditions. But let's also stipulate that those side effects are fairly mild;

³ I use the term "*accessibly could have had*," rather than simply "*could have had*," in recognition of the fact that not all *logically* possible outcomes, however wonderful, are outcomes morality obligates agents to bring about. The fact that it's *logically* possible that one choice leads to a better outcome than another choice doesn't, we think, imply that the other choice is *wrong* in the case where the better outcome is *barred* by, e.g., the laws of physics.

that the child's existence will remain clearly worth having; and that no other treatment nearly as effective as bromocriptine is available.

Those details in place, we, perhaps even more clearly than before, see *nothing wrong* with the woman's choice of c2.

Of course, *prior to choice*, there's no *one* child—p—whom the woman will be assured of conceiving if she conceives at all—it may be p, or q, or r, or any one of many other possible children depending on which sperm accomplishes the actual insemination. After all, there may be 200 million or more sperm cells in a single human ejaculate. (One aside: that biological fact means that the probabilities in Fertility are wildly, upwardly, exaggerated, in order to make the arithmetic we will get to a little later a little more intuitive.) Thus p's chances—indeed any *particular* future person's chances in *any* case—of coming into existence will typically remain low whatever the agent does. But the distribution between wellbeing and chance we see in Fertility applies not just to p but also to *any* child the woman might conceive under the scenario described. And for *any* such child we will think the same thing: that the probabilities at stake in the particular case have converted an otherwise wrong choice into a clearly permissible choice.⁴

2. Expected Value

A widely accepted way of accounting for the intuition of permissibility in cases like Fertility makes use of the concept of *expected wellbeing*, that is, *expected value (EV)*.⁵ Here, by stipulation, it's just the plight of the *one child p* that is at stake; it's how *that* child is, or is expected to be, affected that we need to focus on for purposes here. Now, I think that's how we *should* think about these things; I think, that is, that a person-based approach, rather than an impersonal approach, seems plausible and that we *should* take into account the plights of one person at a time rather than the mass of all persons in aggregate and rather than the *universe* per se. Accordingly, I've constructed the expected value calculation so that we can do just that.⁶

⁴ Prior to choice, the woman—the agent—cannot, of course, refer to any of her potential children by a genuine proper name. But her ability to quantify over classes of individuals for whom she has no such names means that she can nonetheless proceed with an ex ante moral evaluation of her available choices.

⁵ For a clear introduction to expected value theory, see Feldman 2006. In that paper, Feldman argues against expected value theory on grounds of impracticality—grounds other than those I describe in this paper.

⁶ But the formula at work here in the expected value calculation is itself perfectly standard. Thus, if, in the end, we decide *against* thinking about things in that way—one person at a time—we can easily move from what I have written here to an impersonal formula if that's what we want to do.

The *expected value* of a given choice c for a given person i ($EV(c, i)$) =

the summation of the values for i of the outcomes accessible⁷ under c multiplied by the respective probabilities that those outcomes will obtain given c .

We can then calculate as follows:

$$EV(c1, p) = 10 \times .0001 + 0 \times .0099 + 0 \times .99 = .001; \text{ and}$$

$$EV(c2, p) = 0 \times .0001 + 8 \times .0099 + 0 \times .99 = .0792.$$

Since $EV(c2, p) > EV(c1, p)$, $c2$ is the choice that maximizes expected value for p .

It may well seem that—and has indeed in the past seemed to me that—on that basis alone we can then declare $c2$ permissible. Plausibly, the principle that would generate that conclusion for us would be a *simple necessary* condition on wrongdoing, one that says that, when other things are equal as they are in Fertility, a choice that maximizes expected value for p is permissible. Such a principle, of course, leaves open the question whether $c1$ is permissible as well. But it also—plausibly in my view; more on this later—avoids the implication that $c1$ is wrong.

Now, aiming to develop a cogent, plausible, consistent person-based approach, I would add to that simple necessary condition an *existence* condition. The idea would be to make it clear that the principle *doesn't* open the door to the claim (and of course, as a mere *necessary* condition on wrongdoing, wouldn't in any case *imply*) that a given choice is wrong in virtue of the fact that it has failed to maximize expected value on behalf of a person in an outcome in which that person never exists at all.

Putting those two conditions together, we obtain the *expected value maximizing principle (EVMP)*.

Expected value maximizing principle (EVMP). A choice c made in a given outcome x is wrong *only if*

⁷ Here we could just say “possible” rather than “accessible” since any outcome inaccessible under c will be one the probability of which is zero given c . The reverse may not hold: it may be that the probability of a given outcome under a given choice is zero even in a case in which the outcome itself is accessible.

there is a person i who does or will exist in x under c and

there is an alternate available choice c' such that the expected value of c for i is less than the expected value of c' for i .

Under EVMP, c_2 , having *maximized* expected value for p , *isn't* wrong—is, that is, *permissible*. And, for the case of Fertility, that's, of course, the result we want.

Now, some expected value theorists may want to say more—more, that is, than what a simple *necessary* condition on wrongdoing can imply. They may push to say that c_2 is not only permissible but also *obligatory*—that is, that c_1 is *wrong*. In what follows, I will return to that issue, and spell out why I think we may not want to take that further step (part VI below). But our focus, for now, is on what it is about the probabilities that makes c_2 *permissible*. So for now we set aside the issue of whether c_1 is wrong.

3. The Nonidentity Problem

The most challenging nonidentity cases are the ones in which the choice under scrutiny is clearly wrong. Those cases include: the choice of the risky over the safe environmental policy (Parfit 1987); the choice to deplete rather than conserve resources (Parfit 1987); the choice to do nothing rather than something about climate change (Broome 1992); choices yielding historical injustices; and the choice to sign the slave child contract, or to take the iatrogenic pleasure pill, just prior to conceiving a child (Kavka 1981).

Now, it might *seem* that those cases display the same tradeoff between outcome and chance that we see in Fertility. If that's so, and if the argument to permissibility is valid in Fertility, it must be valid in those nonidentity cases as well. And that in turn would mean that, though we *like* the result that EVMP generates in the case of Fertility, we would be compelled nonetheless to reject EVMP on the ground that it generates clearly false results in the nonidentity cases.

But that wouldn't mean that we should abandon the idea that *expected value* has a critical role to play. The defect might instead lie in the EVMP's *existence condition*. We could get rid of that and ask, not whether expected value has been maximized for a *particular* existing or future child p , but rather whether there's *any* child q for whom an alternate choice would have produced still more expected value. If so, then the expected value principle, stripped of the existence condition, would avoid the result that the choice under scrutiny—risky policy, depletion, pleasure pill, etc.—is permissible.

Do we thus have, in the nonidentity cases a compelling objection against EVMP—and, specifically, against the existence condition?

No. The objection relies on the assumption that the pattern of tradeoffs between the better outcome and the better chance in the nonidentity cases is just what we see in Fertility. But that assumption is false.

In Fertility, the woman chooses to undergo a treatment that *increases her fertility* and thus, for her child *p* who in fact exists, made it more likely, *calculated at that time just prior to choice*, that *p*, in due course, would exist. It's just an actual outcome *bias*, an unexamined and faulty *assumption*, an instance of *post hoc ergo propter hoc*, that the choices under scrutiny in the nonidentity cases will do that same work: that choosing depletion over conservation—or choosing the risky over the safe policy, or choosing to do nothing rather than something about climate change, or choosing to take the pleasure pill or sign the slave child contract—will similarly make whoever it is who ultimately exists and suffers as a result of those choices more likely, *calculated at that time just prior to choice*, to exist (Roberts 2007; 2009; Roberts and Wasserman 2016).

And it's an assumption that badly fails when exposed to the clear light of day. The child who is conceived just after the would-be parent takes the pleasure pill is no more likely to exist under that choice than that *same* child is to exist under the choice to take the aspirin instead.⁸

So: so far so good for EVMP, existence condition in place. It generates the correct result for the case of Fertility and avoids clearly incorrect result in the most challenging versions of the nonidentity problem.

4. Two Problem Cases for Expected Value

But we do have, it seems, independent reasons to question both the person-based and the impersonal appeals to expected value.

Consider a second mother and child case, the case of the *All-But-Known Disaster* (Graph 2).

⁸ And it's not just that we can't look into the future and *identify* any such person and say, of that person, that his or her chances of coming into existence are improved. That practical problem is readily resolved in virtue of the fact that we can use quantifiers to talk about those future people. For no such person *i* who might exist following the parent's ingestion of the pleasure pill is *i*'s chances of existence greater under the parent's choice to take the pleasure pill than it is under any alternate, safer choice (e.g., the choice to take the aspirin, or sip of water, in place of the pleasure pill).

Graph 2. All-But-Known Disaster⁹

		.9999	.0001
c1 mom doesn't choose risky extension on behalf of child p EV(p, c1) = 8 x .9999 + 8 x .0001 = 8	+1M +8 +0 -10	o1 p	o2 p
c2 mom chooses risky extension on behalf of child p EV(p, c2) = -10 x .9999 + 1M x .0001 = 90.001	+1M +8 +0 -10	o3 p	o4 p

$M = 1,000,000$.

Two important features of this case should be highlighted. First, in this new case, the child p faces *no* existential risk. Perhaps p already exists or that p will exist whatever choice is made. Second, c2 is a single, one-off, high risk/high reward, choice. To be clear: the option of the mother's, the agent's, pursuing a *long-term* strategy that allows for the choice of c2 *in the context of* a sequence of choices that includes not just c2 but also successive c2-like choices, choices that repeat the risks and rewards generated by c2, is *not* a choice available to the agent. Thus, p may end up—and in the particular case *does* end up—as the *single trial non-beneficiary* of the high risk/high reward c2.

To give some color to the case, we can suppose that a mother faces the choice whether to have her only child, p, undergo a very risky but possibly enormously beneficial treatment, a treatment that might extend p's life for thousands of years. If she declines treatment—chooses, that is, c1—then p is certain to have, at +8, a good life however the future unfolds. Specifically, let's imagine that at +8 p's wellbeing level is about what the woman herself, as well as what others in her and her child's generational cohorts, will enjoy. In contrast, if she chooses treatment—chooses, that is, c2—then the probability is *very high* that p will have, at -10, an existence that is significantly *less* than an existence worth having. In fact, she *all but knows* that, if she chooses c2, p will be thoroughly miserable; she *all but knows* that c2 will be a complete disaster for her child. And it is.

Is her choice of c2 wrong? We would easily, unobstructedly, say that it is wrong, *but for one final detail of the case*: the very small probability that the choice of treat-

⁹ I am grateful to Dean Spears for the case on which All-But-Known Disaster is based.

ment will unfold into an outcome that benefits *p* enormously—that generates for *p* an extraordinarily good life, say, a life that is *both* very good *and* goes on and on for thousands of years.

Here, we calculate as follows:

$$EV(c1, p) = 8 \times .9999 + 8 \times .0001 = 7.9992 + .0008 = 8; \text{ and}$$

$$EV(c2, p) = -10 \times .9999 + 1M \times .0001 = -9.999 + 100 = 90.001.$$

If the concept of expected value is the right way to bring probability to bear in moral analysis—if a choice that maximizes expected value for one person and leaves everyone else alone can't be wrong—then we should conclude that *c2* is permissible.¹⁰

But can the result that *c2* is permissible be correct? Do we think, in the case of All-But-Known Disaster, that the fact that *p* might—*might; against all odds*, the chances of things turning out well for *p* under *c2* being a scant 1 in 10,000—obtain an enormous benefit converts what would otherwise be a clearly *wrong* choice into a *permissible* choice?

I don't think that we do. The concept of expected value helped to make sense of our intuitions in Fertility *and* the nonidentity cases. But it *doesn't* help us make sense of All-But-Known Disaster *at all*. Though *c2* clearly maximizes *expected value* for the one existing or future person *p* who might possibly be affected by how the choice is made, our clear intuition is that *c2* is *wrong*.

To test that intuition further—to make the case more vivid, but not to change facts of the case or in the end the moral analysis—let's just note that it's part of the case that the woman *in fact* chooses *c2* and that the predicted outcome, *o3*, the outcome in which disaster strikes, *in fact* obtains. We see the miserable child; we see the hapless mother. It now seems clearer than ever that *c2* is wrong.

Of course, in *another case altogether*—call it *Long Run Risk Reduction*—we may well say that *c2* is permissible. In that case, the woman has the option of completing a *sequence* of high risk/high reward choices, commencing with *c2* and continuing with (perhaps thousands) of successive *c2*-like choices, choices that repeat the risks and rewards for *p* that come with *c2* itself. *Provided* that the available sequence is sufficiently long *and* that the mother in fact plays out the long-run strategy—provided, that is, that *c2* is chosen in (relative to) an outcome that meets those con-

¹⁰ Indeed at least some theorists want to go beyond the result that *c2* is permissible and say more than a simple necessary condition on wrongdoing can imply—that is, that *c1* is *wrong*. But again for the moment we'll set that issue aside and focus exclusively on the question of *c2*'s permissibility.

ditions—c2 plausibly *is* permissible. For then the long term strategy, itself commencing with c2, will come with a significantly high probability that *at some point* p will luck out and accrue an enormous benefit, thereby, at least over the long run, significantly reducing the risk of disaster for p.

But All-But-Known Disaster is *not* that sort of case. In All-But-Known Disaster, the accessible outcomes that include the choice of c2 are limited to just o3 and o4. And in each of those outcomes it's stipulated that c2 exists as just a *single, lonely, one-off* choice and *not* as the start of a risk-reducing sequence of further c2-like choices.¹¹

One other quick note. Some theorists might bring the concept of risk aversion into play at this point to explain away the intuition that c2 is clearly wrong, or perhaps to allow us to say that, relative to the risk averse agent, c2 is wrong, but relative to the risk neutral agent, c2 is permissible. And it's not implausible to suppose that those claims can form the start of a plausible account of Long Run Risk Reduction.

But All-But-Known Disaster is a very different sort of case. In that case, it seems highly implausible that even the evaluator *neutral* in respect of risk will consider c2 permissible. Even the *neutral* evaluator will agree that it's wrong for the woman, when she *all but knows* c2 will end in disaster for her child, to choose c2. Considerations of risk aversion thus seem beside the point in All-But-Known Disaster.

Here is a second problem case for the expected value approach, one that does involve existential risk.

¹¹ In general, it's a mistake to think that what we say about the permissibility of a choice (an act) made (performed) in one outcome in one case must always be what we say about the permissibility of that same choice (act) made (performed) in another outcome in another case. Thus the fact that c2 is permissible relative to (when chosen in) certain outcomes available to the agent in Long Run Risk Reduction gives us *no reason at all* to say that c2 isn't wrong relative to (when chosen in) o3 in All-But-Known-Disaster. At least, that that would be a mistake is an essential tenet of any consequentialist approach. As Feldman points out, the same medicine might either save or kill a patient, depending on whether or not it's followed up by a second medicine, and we can consistently take the position that giving the one medicine is permissible relative to (when made in) the outcome in which it's followed up by the second but impermissible relative to the outcome in which it's not. Feldman 1986.

Graph 3. 50-50 Disaster¹²

		.5	.5
c1 agent chooses to bring child into existence under high risk/high reward condition	+101	<i>o1</i>	<i>o2</i>
	+0	p	
	-100		p
c2 agent chooses not to bring child into existence at all	+100	<i>o3</i>	<i>o4</i>
	+0	<i>p</i>	<i>p</i>
	-100		

The potential upside for *p* being just a bit greater than the potential downside, $EV(c1, p) > EV(c2, p)$. According to EVMP, *c1* is permissible. But that seems implausible. Where the agent has the option of not bringing the child into existence to begin with—the option to assure that the child will *not* be forced to endure a deeply, genuinely *wrongful* life—isn't the only permissible choice the choice not to bring the child into existence to begin with?

* * *

Thus the appeal to expected value turns out to be problematic. Yet we can't let go of the idea that probabilities play a critical role in moral evaluation. Fertility and many other cases as well tell us that that's so. My own proposed solution to the nonidentity problem in its most challenging forms, moreover, depends critically on that being the case.

The question thus arises whether probabilities can be brought to bear in the analysis by a route other than that of expected value.

5. Probable Value

One such possible route lies in the concept of *probable value* (PV). The underlying idea is that what makes the difference between the moral status of *c1* and *c2* in All-

¹² I am grateful to Mark Budolfson for this case. Risk and Population Workshop, University of Texas, Austin, Texas (Nov. 22-24, 2019).

But-Known Disaster is a substantially narrower set of facts than what we are required to take into account in calculating expected value. The EV calculation takes *all outliers* into account. In All-But-Known Disaster, the tiny chance of the extraordinary o4 obtaining under c2 forced us to say that c2 was permissible.

Under a probable value approach, the tiny chance of o4 obtaining would not be relevant to the mechanics of bringing probability to bear in the context of moral evaluation. Rather, under a probable value approach, where c2 ends in disaster, as in o3, what is critical in evaluating c2 at o3 is just the high probability, calculated just prior to choice, that o3 would in fact unfold under c2. The very high probability that o3 will obtain given c2 is the probability-significant feature that we are to refer to when evaluating c2 relative to o3. It's *that* number that we should multiply by the value that o3 has for p, to produce the *probable value of o3 for p under c2* (PV(o3, p, c2)). And the necessary condition on wrongdoing that we should adopt is the failure to maximize not *expected* value but rather the failure to maximize *probable* value.

The definition of probable values comes in two parts:

Probable Value.

Where a choice c made in (relative to) an outcome x creates a probability n (calculated just prior to choice on the basis of information available to agents at that time) that a person i will have the wellbeing level (wb) that i in fact has at x, the *probable value* of x for i under c = the wellbeing i has at x multiplied by n.

Where the *minimal wellbeing level* (mwb) for a person i at the range r of outcomes that may accessibly obtain under c is the *least* wellbeing level i has in any member of r, and where that choice c creates the probability n (calculated just prior to choice on the basis of information available to agents at that time) that some member or another of r will obtain, the *probable value* at r for i under c = n(mwb).

That is: $PV(x, i, c) = n(wb \text{ of } i \text{ at } x)$; and, where the probabilities of various outcomes under a given alternate choice are each very low but the alternate choice nonetheless *dominates* (we can say) the one choice, $PV(r, i, c) = n(mwb \text{ of } i \text{ at } r)$.¹³

In All-But-Known Disaster, then, the probability given c2 that p will end up at a wellbeing level of -10 is .9999, making $PV(o3, p, c2) = -9.999$. In contrast, we see

¹³ I am grateful to Tomi Francis for a case that demonstrates the need for the second part of the definition of probable value. Nonidentity Workshop, Institute for Future Studies, Stockholm (Feb. 8-9, 2020).

higher probable values for both c1 at o1 and c1 at o2: $PV(o1, p, c1) = 8 \times .9999$; and $PV(o2, p, c1) = 8 \times .0001$.

The simple necessary condition, then, on wrongdoing would be that a wrong choice can't be one that maximizes probable value for each and every existing or possible person. For the person-based take on things, we would, as before, add the existence condition. The result is the *probable value maximizing principle (PVMP)*:

Probable value maximizing principle (PVMP). A choice c made at (relative to) a given outcome x is wrong *only if*

there is an individual person i who does or will exist in x and there is

(i) an alternate available choice c' made at (relative to) an alternate accessible outcome y such that $PV(x, i, c) < PV(y, i, c')$ or

(ii) z is a member of the range r of outcomes that may accessibly obtain under c' and $PV(x, i, c) < PV(r, i, c')$.

There are, then, three ways of satisfying the proposed probability-related necessary condition on wrongdoing: one is for there to be an alternate choice at an alternative outcome that creates more probable value; another is for there to be an alternate choice and an alternate range of outcomes that shows that the alternate choice dominates the one choice; and still another is for the person for whom things look to be mathematically unfortunate never to have existed at all in the particular outcome at all.

Since $PV(o3, p, c2)$ *isn't* maximized—it's *less* than $PV(o1, p, c1)$ and *less* than $PV(o2, p, c1)$ —nor is it dominated by any alternate choice, PVMP nicely avoids the false result that c2 at o3 is permissible.

Thus PVMP—in contrast to EVMP—leaves the door open for us to say (under still other principles) that c2 at o3 is wrong. *That is progress.*

6. Actual Value Condition on Wrongdoing

Let's go back to Fertility. I think we can mine that case for still a third necessary condition on wrongdoing. Consider c1 at o1. There, p, against all odds, has made it into existence even though the woman *hasn't* chosen to undergo the fertility treatment, and p's maximized wellbeing has come at no cost to anyone else. It seems plausible to me to say that c1 made at (relative to) o1 is permissible.

The principle that says just that is the *actual value maximizing principle (AVMP)*.

Actual value maximizing principle (AVMP): A choice made at (relative to) a given outcome is wrong *only if* there is a person who does or will exist in that outcome and an alternate available choice and an alternate accessible outcome such that the actual value of the one outcome for that person under the one choice is less than the actual value of the alternate outcome for that person under the alternate choice.

In other words: AV maximization for a given person, other things equal, implies permissibility.

7. Note on the Existence Condition

The existence condition, as things stand, has been built into both PVMP and AVMP. But I want to underline that the condition *isn't* that the person exist under *both* choices or in *both* outcomes. The requirement is just that a person for whom things are worse does or will exist in the one outcome.

In other words: that a person never exists at all, other things equal, implies permissibility.

The Narvesonian idea behind the existence condition is that the fact that a person never exists in a given outcome under a given choice doesn't, morally, count against that outcome or against that choice: that a person *doesn't* exist doesn't, other things equal, make the outcome worse or an otherwise permissible choice wrong. Under that principle—the existence condition—we find in Fertility that c2 at o3 is perfectly permissible. For the same reason, so is c1 at o2.

8. Consolidated Principle

We can consolidate the three various necessary conditions noted above into a single principle, one that makes use of the concept of probable value in explaining how probabilities are to be brought to bear in our evaluation of choices; that accepts that a person's *nonexistence* doesn't (other things equal) count either against the outcome or against the choice; and that accepts that, when we (whether against all odds or not) *succeed* in maximizing wellbeing for a given person, what we have done isn't, in the end, wrong (though it may well have *looked* wrong starting out).

I should note that the combined principle does more than just sum up the three necessary conditions. It also takes the substantive position that variations in proba-

bility bear solely on the evaluation of the agent's alternative *choices* and not at all on the ranking of *outcomes* in respect of their overall betterness. That means that, while embedding in an outcome a probability of—for example—.0099 in place of a probability of .0001 may be critical to the evaluation of the choices made within that outcome, it *isn't* critical to the ranking of the one outcome against still other accessible outcomes; it can't, on its own, make one outcome better or worse than another. What makes one *outcome* better or worse than another is a matter confined to how wellbeing is distributed across a given population and what alternate outcomes are accessible in the particular case. What makes one *choice* permissible rather than wrong may turn on the probabilities.

Thus the combined principle must be spelled out in two parts, the telic (relevant to the ranking of outcomes) and the deontic (relevant to the evaluation of choice).

PVMP+AVMP+EC:

Telic component: Where y is accessible relative to x, x is worse than y, only if there is a person i and an alternate outcome z accessible relative to x such that:

i does or will exist in x; and

x is worse for i than z is (where z may, but not need, be identical to y).

Deontic component: A choice c performed at x is wrong, only if there is a person i, an alternate choice c' and an alternate accessible outcome y accessible relative to x such that:

i does or will exist in x;

x is worse for i than y; and either

$PV(x, i, c) < PV(y, i, c')$ or y is a member of the range of outcomes that may accessibly obtain under c' and $PV(x, i, c) < PV(r, i, c')$.

On the choice side, this principle provides a complete evaluation of Fertility: all choices at all outcomes are permissible (yay!) and a nearly complete evaluation of All-But-Known Disaster: c1 at o1 is permissible, and so is c2 at o4. But what does it say about the nonidentity problem?

9. The Nonidentity Problem Revisited

Earlier, I claimed that EVMP opens the door to a plausible account of the most challenging nonidentity cases.

In this part, we turn to whether PVMP provides an equally plausible account of the nonidentity cases.

And I'll just note that it easily does. Let's, just to keep things simple, focus on the pleasure pill case. Since the burdened child p's chances of existence are just as great whether the parent takes the pleasure pill or—say—an aspirin, the choice to take the pleasure pill relative to the outcome that in fact unfolds—the outcome in which the child in fact exists and suffers from the pill's side effects—fails to maximize PV. Nor does the choice maximize AV. And, finally, we simply note that the burdened child in fact exists in the particular outcome and under the particular choice under scrutiny.

That means that all three necessary conditions identified in the consolidated principle are satisfied. That fact, in turn, opens the door for still other principles—other person-based principles, I would surmise—to step in and say that the choice that the parent has made at that outcome is wrong.

To put the point another way: the concept of probable value positions us to provide a plausible account of the pleasure pill case and other nonidentity cases just as nicely as the concept of expected value does—which is to say, quite nicely.

10. A Problem for Probable Value: The All-But-Known Success and the Unexpected Really Bad Outcome

At the same time, issues remain. Consider a new case, All-But-Known *Success*, a classic case that seems to cry out for the concept of expected value. Obviously, however, adding still a fourth necessary condition rooted in expected value won't work since we would then get false results in All-But-Known Disaster. So what are we to do with cases like All-But-Known *Success*? In this new case, the probabilities are clearly relevant. But by what principle are they to be brought to bear?

Graph 4. All-But-Known Success

		0.9999	0.0001
c1 undertakes mission	+10 +1 +0 -1000	<i>o1</i> p PV: 9.999	<i>o2</i> p PV: -0.1
c2 doesn't undertake mission	+10 +1 +0 -1000	<i>o3</i> p PV: 0.9999	<i>o4</i> p PV: 0.0001

Again, to add color to the case, let's suppose still another mother-and-child scenario—and another one-off case. Here, the woman must choose whether to undertake a mission to rescue her only child *p* from space aliens who have confined *p* in a cage allowing *p* only the minimal conveniences of life (nutrition, hydration and occasional ablution) or to do nothing and leave *p* to live out a life only *barely* worth living. If the woman's mission succeeds, *p* will be restored to a wellbeing level of approximately what the woman herself will have and what others in her and her child's generational cohorts will have. Because of the woman's overall proficiency and vast experience in previous extremely dangerous search-and-rescue efforts, the chances that her mission will succeed is extremely high. On the other hand, if her mission fails, the space aliens are sure to become aware of her efforts (her transport vehicle leaving a signature trace in the alien atmosphere, easily recorded by their sophisticated and highly sensitive monitoring devices) and they will (certainly) respond by torturing *p* in horrible ways for the remainder of *p*'s very long natural life, leaving *p* with an existence far less than one worth having.

Let's suppose, too, that the woman chooses to proceed with the mission—chooses, that is, *c1*. And let's suppose as well that, against all odds, the mission fails, and *p* is left to suffer in the horrible ways we have just described.

This is the sort of scenario that reminds us that to *all but know* something is not to *know* something.

The question is: both AV and PV being relatively low, and the existence condition being satisfied, and EV being off the table, on what grounds (if any) do we say that

the woman's choice of c1 is permissible *even though* it has ended in disaster for p?

As already underlined, the consolidated principle PVMP+AVMP+EC, offering mere *necessary* conditions on wrongdoing, has no capacity, for any choice, to generate the result that that choice is wrong. It can at most generate, for a given choice, that that choice is permissible.

That means that the problem for the combination principle isn't that it generates a *false* result—the result that c1 is wrong—but rather that it's *incomplete*: that it *doesn't* generate the result that c1 is permissible.

It's—again—the probabilities that are at stake that make us want to say that c1 is permissible—that in All-But-Known Success, the mother all-but-knew, prior to choice, that things would work out well for p. But by what principle are the probabilities brought to bear to determine that c1 at o2, where disaster has struck, is also permissible? We can't appeal to the fact that PV or AV is maximized or to p's nonexistence, and EV is off the table. But what else is there?

11. Two options: Appeal to epistemic security or accept choice as wrong

To address this issue, we seem to have two options. The first exploits the fact that the most distinctive feature of All-But-Known Success is the *very* high probability of success (.9999) in that case in combination with the fact that p has a lot at stake (will p have a life well worth living, or a life just barely worth living?).

But what's the principle?

It seems that we can, in advance of a principle that we can actually test, say some things that seem both relevant and correct. We can say that, when the agent's position is *epistemically secure* relative to the better outcome and the agent all-but-knows that things will turn out well, as in All-But-Known Success, then the permissibility that comes with a choice that ends in an outcome that *does* turn out well can be *attributed* to the choice that ends in an outcome that (against all odds) *doesn't* turn out well. In contrast, when the agent's position is *epistemically weak* relative to the successful outcome, any such attribution of permissibility would be inappropriate. For in that sort of case, the agent can't plausibly be said to all-but-know, much less *know*, that things will turn out well. Ditto All-But-Known Disaster.

This isn't, of course, a principle we can test; it's just a bare-boned description of a schematic for such a test, one that links moral evaluation not just to the usual suspects, that is, facts about outcomes and the probabilities that those outcomes will obtain under the choices under scrutiny, but also to the agent's particular epistemic state. From there, the question remains whether an appropriate fourth necessary condition on wrongdoing can be fashioned.

The second is to take the position that we aren't challenged, in this case, to find grounds for the claim that the choice that ends in disaster—here, c2 at o2—is permissible. We should be comfortable with the fact that the door is open to a finding of wrongdoing in this particular case. For perhaps in this particular case the choice (though it *looked* permissible starting out) really *was* wrong.

References

Broome, John (1992). *Counting the Costs of Global Warming*. Cambridge: White Horse.

Kavka, Gregory (1981). "The Paradox of Future Individuals." *Philosophy & Public Affairs* 11: 93–112.

Parfit, Derek (1987). *Reasons and Persons*. Oxford University Press (originally published 1984).

Roberts, M.A. (2007). "The Nonidentity Fallacy: Harm, Probability and Another Look at Parfit's Depletion Example." *Utilitas* 19: 267–311.

Roberts, M.A. (2009). "The Nonidentity Problem and the Two Envelope Problem." In *Harming Future Persons*, eds. M.A. Roberts and D. Wasserman. Springer, pp. 201–228.

Roberts, M.A. and David Wasserman (2016) "Dividing and Conquering the Nonidentity Problem." In *Current Controversies in Bioethics*, eds. Matthew Liao and Collin O'Neil. Routledge, pp. 81–98.