

## Does the Moral Significance of Merely Possible People Imply that Early Abortion Is Wrong?

Melinda A. Roberts

The College of New Jersey

[robertsm@tcnj.edu](mailto:robertsm@tcnj.edu)

*Abstract.* Do merely possible people—people who, relative to a given possible future or world, *never* exist—*matter morally*? I begin this paper by showing that the position that they don't is completely untenable. But the position that they do seems to open the door to a powerful argument against early abortion—an argument that goes through *even if* it also happens to be the case that the human embryo and early human fetus, as non-thinking things, have no moral status at all in their own right. The purpose of this paper is to reconcile the positions that (1) the merely possible matter morally and (2) leaving any such person out of an existence worth having imposes a *loss* on that person (is a way, that is, of creating less wellbeing for a person when agents could have created more) against the position that (3) the early abortion is often perfectly permissible. *Variabilism*, I argue, accomplishes that reconciliation. According to Variabilism, each person—you, me and the merely possible—matters *morally* but we each matter *variably*. That is: some of our losses have full moral significance and some have no moral significance whatsoever, with moral significance itself being a function of *where* the loss is incurred in relation to the person who incurs it. Specifically: the loss a person incurs at any world has moral significance just in case that loss is incurred at a world where that person does or will exist; such a loss otherwise has no moral significance whatsoever. The point that Variabilism establishes is a simple one. It allows us to take the position that the early abortion, whether chosen for a good reason, a poor reason or no reason at all, often remains perfectly permissible, but also to recognize both that the person the early abortion relegates to the class of the merely possible matters morally and that leaving that person or any other out of an existence worth having is indeed to create less wellbeing for that person when we could have created more.

### 1. *An Argument Against Early Abortion*

The purpose of this paper is to show that we can take the position that merely possible people matter morally in just the same way you and I matter morally, and also the position that leaving a person out of an existence worth having is to make things worse for that person when we could have made them better, but at the same time (easily) take the position that the early abortion is often permissible, whether it is chosen for a good reason, a bad reason or no reason at all.

It may not be immediately obvious that the first two positions that I have described here are claims we should accept. On reflection, however, as we shall see, they are hard to resist.

1. *The Moral Significance of Merely Possible People (MSMPP)*: People who will never exist at all matter morally in just the same way that people who do or will exist matter morally. Making things better for the merely possible rather than worse—creating more wellbeing for them rather than less—is just as important, morally

speaking, as making things better rather than worse for existing and future people is. What we might think of as the basic maximizing insight—that we ought to do the best we can; that we ought create the most good for people that we can—applies to *them* in exactly the same way that it applies to *us*.

2. *The Betterness (Worseness) of Existence (Comparability)*: In many cases, it is *better for a person* that the future has unfolded in a way that includes that person rather than in a way that excludes that person. The *actual* world, for example, in which *we* exist, is *better for us* than is any alternate possible world in which we never exist; we have *more wellbeing* at the actual world, where we exist, than we have at an alternate world where we never do. Consistent with that point, there may also be extraordinary cases in which it is *worse for a person* that the future has unfolded in a way that includes that person rather than in a way that excludes that person. If the *actual* world, for example, includes any such people, then the actual world is *worse for each such person* than is any alternate possible world in which that person never exists at all.

MSMPP, in particular, may seem less than obvious—at least to anyone who has ever thought there might be *something* to Narveson's idea that our obligations have to do with "making people happy" rather than "making happy people."<sup>1</sup> Why, after all, should people who *never exist* have exactly the same moral status that you and I have? I don't know why they *should* but the fact is—as we shall see—they *do*. We are going to have no choice in the matter; we shall have to accept MSMPP.

Comparability raises controversies of its own. At least, it may seem odd to say that existing can be better (or worse) for a person than never existing at all. And some philosophers have actually sketched out arguments against Comparability. Such arguments often take the following form: if existing is *better* for a person than never existing, then never existing would have been *worse* for that person. But in that latter case precisely *who* would never have existing have been worse *for*? Surely, no one, in which case never existing cannot, after all, be worse for that person.<sup>2</sup>

There is, however, an intuitive objection to this argument. We can avoid the charge of incoherence by simply answering the question: the person never existing would have been worse for is *the person who has the existence worth having*. That is something speakers seem able cogently to say, whether the world of utterance is the world at which that person exists and

is happy, or alternatively the world at which that person never exists. At the one world speakers can say: existence is better for *this person*, and it is for *this person* that never existing would have been worse. That is: *this person* has more wellbeing at this world where he or she has such-and-such a happy existence than *this person* has at any alternate world where he or she never exists at all. And at the other speakers can say: existence is better for any person who has such-and-such a happy existence, and for any such person, never existing would have been worse than having such-and-such a happy existence. That is: any such person has more wellbeing at a world where that person has such-and-such a happy existence than that person is at any alternate world where he or she never exists at all.<sup>3</sup>

If reference is not the issue, perhaps the obstacle is the comparison itself. But we can surely imagine that there will come a day—in, of course, the *very* distant future—in which we think that additional moments of life will not make things better for us and may even make things worse for us. Once reference is settled, it is unclear we cannot make the same comparison for the whole life rather than just a portion of that life.

So: we have two claims that are going to be hard to resist. MSMPP we are—as we shall see—compelled to accept. Comparability is at least hard *enough* to resist that we wouldn't want any arguments or positions that we develop here or elsewhere *to count on* its being false or nonsensical. In particular, a defense of early abortion that holds up *only if* Comparability at the end of the day turns out to be (no doubt on highly technical grounds) false would strike me as tenuous. Similarly, David Heyd's argument that it is permissible to leave the additional happy child into existence—his defense, that is, of the second half of what McMahan called *the Asymmetry*—strikes me as tenuous for that very reason.<sup>4</sup> If our *only* way to obtain that highly intuitive second half is to reject, with Heyd, Comparability, I would at least lose confidence in that particular intuition.

Yet MSMPP and Comparability together at least appear to throw into question the permissibility of early abortion in any ordinary case. Let's take the simplest case first. Let's suppose that the early abortion makes things better for no one at all. And let's suppose as well that the early abortion is worse for the possible future person—that is, that the existence itself is worth having. According to MSMPP, the possible future person the early abortion would cause never to exist at all *matters morally*, and in exactly the same way that you and I matter morally. It is just as important to make things better for possible persons as it is to make things better for existing and future persons. Comparability, moreover, tells us that bringing a person into existence can make things better for that person: it can be better—and often is better—to exist than never to exist at all. We then seem compelled to conclude that the early abortion is wrong.

Moreover, this argument seems to go through, even if we take for granted that the human embryo and early human fetus themselves have no moral status whatsoever—that is, that they themselves are not the kinds of things in respect of which we are generally obligated to create more wellbeing for rather than less. For on this way of looking at things, the early abortion is wrong, not in virtue of what it does to the human embryo or early human fetus, but rather in virtue of what it does to the possible future person—in virtue, that is, of the fact that it leaves that person out of existence altogether, thereby creating less wellbeing for an individual who has moral status—a *person*—when agents could have created more at no cost to anyone else at all.

Of course, in many cases, the early abortion will make things not just *worse* for the person the early abortion causes never to exist at all but also *better* for one or more people—perhaps the pregnant woman herself, her partner, her already-existing or future offspring or still others. In many cases, in other words, the facts can be complicated. Still, the early abortion will very often come out wrong. The difficulty is that the depth and dimension of the loss for the possible future person, the person who never exists if the early abortion is performed, is vast;

it's everything. It is so grave that, *even if* the early abortion is better for others, the early abortion is still wrong. We cannot justify imposing a very deep and great loss on one person in order to avoid a relatively shallow loss on behalf of still another. And that is so, however we think conflict, or tradeoff, scenarios are to be resolved—whether by comparing the dimensions or relative depth of the competing losses or by aggregating the competing gains.

Accordingly—taking into account both the simple case (no tradeoffs) and the more complicated case (some tradeoffs)—the implication from MSMPP and Comparability may well seem to be just this:

3. *Conclusion Against Early Abortion.* The early abortion is often wrong. The purpose of this present paper is to contest the inference to (3). Consistent with MSMPP and Comparability, we can recognize that the early abortion in the ordinary case is perfectly permissibility, whether chosen for a good reason, a bad reason or no reason at all.

\* \* \*

The particular power of this argument—henceforth, the *MSMPP Argument Against Early Abortion*—is in the fact that, if it works at all, it works *even if* the human embryo and early human fetus have no moral status whatsoever. As noted before, the argument isn't about *them*; it's about the *possible future person*.

One might have thought that establishing that the human embryo and early human fetus lack moral status—that they are not the kinds of things in respect of which we have moral obligations for their own sake—would itself be sufficient to establish in any ordinary case the permissibility of the early abortion. It might have seemed, moreover, that two very plausible positions regarding the kinds of things that matter morally and the neurology of the human embryo and early human fetus would together establish that very lack of moral status. Those positions include, first, the *Thinking Thing Account of Moral Status*.

*The Thinking Thing Account of Persons.* In determining what we ought to do, the kinds of things whose interests we must take into account—the kinds of things that

matter morally *at least if they do or will exist*—are the kinds of things that feel, understand, learn or anticipate—in short, the kinds of things that think. Such things come into existence upon their first thought and go out of existence at the conclusion of their last. For purposes here, we call such things *persons*.<sup>5</sup>

Why would thinking (in this broad sense) be a *necessary condition* for mattering morally?

Perhaps it is just that it is *only* in the case of thinking things that it *matters to them* (whether they think so or not) that we treat them in one way rather than another.

We then note as a matter of empirical fact that:

*Human Embryo and Early Human Fetus as Non-Persons.* The human embryo and early human fetus *do not think*—they do not feel, understand, learn or anticipate—and so are not *persons* in that sense.

Consistent with this point, we can also (plausibly) say that the *late* human fetus overlaps spatially for a period of time with a *person*; we might even say that during that brief period the late fetus itself *thinks*, implying, not that the fetus *is* a person but rather that there is a person “in the vicinity” of the fetus. But the human embryo or early human fetus, never having had a thought, clearly does not think. It is clearly not a person—as I use that term here—nor does it signify a person “in the vicinity.” Moreover, even that very first thought may be irrelevant from the moral point of view unless it can connect with future thoughts in such a way as to allow the person who has come into being to *survive* in the absence of the early abortion. That is: if the person was due to go out of existence as a matter of natural law whether the early abortion is performed or not, then the early abortion could not itself make things worse for that person.

For clarity, we can now just specify that the *early abortion* is the abortion that takes place *prior* to the point at which the fetus has developed into a person, a thinking thing. Before we met MSMPP, we might have thought that the Thinking Thing Account was sufficient to ground the permissibility of the early abortion. We might have thought that, on the Thinking Thing Account, the early abortion would be perfectly permissible in any case in which the *only* person the early abortion makes things worse for is the person the early abortion causes never to exist at all. And we might have thought as well that the early abortion would often be perfectly

permissible even if the early abortion *does* make things worse for others beyond the possible future person. In such cases, the permissibility of the early abortion would then depend on how the various conflicts are resolved—with it being understood that, ordinarily, the conflict between the pregnant woman who wants the early abortion and anyone else, e.g., the *grandmother* who wants a great-grandchild, is to be resolved in favor of the pregnant woman.

But we then meet MSMPP and determine—as we shall see—that we must accept MSMPP. We see that the permissibility of the early abortion depends on more than a showing that the human embryo and early human fetus are not *persons*, that is, are not the kinds of things in respect of which we have moral obligations for their own sake. We see that the needs and interests of the possible future person must become part of the mix as well. Suddenly, the issue of early abortion, and indeed nonconception, becomes highly fraught—and it begins to seem that the permissibility of the early abortion, and indeed the abortion that makes things worse for *no one* other than the possible future person, is extremely uncertain and indeed tenuous at best.

The goal of this paper is to show what is wrong with the MSMPP Argument Against Early Abortion. Specifically I explain how MSMPP can be true but fail to support in even the smallest way the conclusion that early abortion is wrong. Thus I concede, and believe, that MSMPP is true. And, as MSMPP instructs, it won't do to say that "we" existing and future people matter morally and that "they"—all those people who will never exist at all; all those *merely possible* people—matter not at all. I will argue, however, that it doesn't follow that *every* occasion on which agents create less wellbeing for you, me or the merely possible when they could have created more—that *every* downward shift, *every* diminution, in wellbeing that any of us is the subject of—*counts against* what has been done. Nor, I will argue, do the various problem cases suggest that the fact that we—and they—matter morally comes to anything more than just that:

that you and I matter morally, and that any merely possible person matters morally, comes to just the point that sometimes what is done to us matters morally and sometimes it does not.

The terminology of “loss” is efficient for the purpose of making this point—and, in the end, for the purpose of articulating just when what is done to a person matters morally and when it does not. As I use the term here, “loss” is *shorthand* for any case in which agents have created less wellbeing for a person when they could have created more wellbeing for that same person—any case, that is, in which agents have made things worse for a person when they could have made them better. The defect in the MSMPP Argument Against Early Abortion lies in the following line of inference: if (A) *some* of a given person’s losses matter morally—and the cases show that they do—then (B) that person matters morally; and, if (B) the person matters morally, then (C) *all* that person’s losses matter morally. I am perfectly happy to grant (B) on the basis of (A); (B) is, at least, a legitimate way of making an extremely important point. But we have no basis for granting (C) on the basis of (B). For one thing, we have no basis for thinking that (C) follows from (B). Unless (B) goes far beyond (A), (C) clearly does not follow from (B). Yes, to say that a person matters morally is to say that *some* of that person’s losses matter morally. We can’t explain the cases without conceding that much. But there is nothing in the claim that *some* of that person’s losses matter morally that rules out that still *other* of that same person’s losses matter not at all. For another, we have every reason to accept what the cases themselves seem to suggest—that (B) is true but that (C) is false.

*We do* exert a certain moral pull over the acts of other agents whether those acts are performed at this world, where we do exist, or at worlds where we never exist at all. The *merely possible* exert that same moral pull over our own acts, both at this world, where they never exist, and at still others, where they do exist. That’s MSMPP. But that cross-world moral pull, I propose, is *variable* in nature rather than *constant*. With respect to losses of one sort, that pull

is in full force, whatever its subject's *modal* status. With respect to losses of the other sort, there is no pull at all, even if the subject of that loss happens to be you or me.

## 2. Hare's "Golden Rule" Argument Against Abortion

A version of the MSMPP Argument Against Early Abortion appears in the paper by R.M. Hare. Hare crafts the argument in terms of what he calls the "Golden Rule," which, he says, states that "we should do to others as we wish them to do to us."<sup>6</sup> He then wants us to infer the following: "[i]f we are glad that nobody terminated the pregnancy that resulted in our birth, then we are enjoined not, *ceteris paribus*, to terminate any pregnancy which will result in the birth of a person having a life like ours."<sup>7</sup> He endorses Comparability—and flatly rejects the "assumption" that "one cannot harm a person by preventing him coming into existence."<sup>8</sup> And he argues against what might seem a commonsense antidote to MSMPP—the idea that the "potential people" he enjoins us to bring into existence lack the "right or interests" we should ascribe only to "actual people." For one thing—and here he possibly anticipates the Rabinowicz principle<sup>9</sup>—he thinks it "strange" to say that "by preventing the existence of the object of the wrongdoing, it would remove its wrongness. This is too easy a way of avoiding crime."<sup>10</sup> For another, he notes that "*present* nonactuality" doesn't make a moral difference to us: we recognize the interests of people who do not yet but will exist.<sup>11</sup> So why not—he suggests—recognize as well the interests of people who could but never do exist? Along the way, he remarks on certain connections between the Golden Rule argument against early abortion and the "classical or total utility principle," which he, on other grounds, also accepts. For him, the total principle—what I will call *Totalism*—imposes certain limits on the stringency of the procreative obligation that the Golden Rule argument itself burdens us with: early abortion becomes perfectly permissible only at that point at which not having the early abortion, that is, bringing the pregnancy to term, fails to increase aggregate wellbeing.<sup>12</sup> Totalism, moreover,

reinforces the Golden Rule argument: the procreative obligation we obtain from the Golden Rule, as limited by Totalism, is precisely the procreative obligation we obtain from Totalism itself.<sup>13</sup>

So here we are, back in 1975. What relevance does Hare's argument have for the debate of the day—the debate regarding the moral significance of merely possible persons?

The point is just this: the most seemingly intuitive, commonsensical, “I've got *my* feet on the ground” objection to Hare's argument just won't work.

The objection that *doesn't* work can be put this way. Hare is effectively begging the question against early abortion by assuming that the possible future person the early abortion keeps from ever existing at all has just the same moral status that you and I have. He is begging the question against early abortion, in other words, by assuming that that possible future person is protected by the Golden Rule in the same way that you are protected by the Golden Rule in your dealings with me and vice versa. The better view, according to this objection that *doesn't* work, is this: the class of people protected by the Golden Rule and indeed the class of people who matter morally is a great deal *more exclusive* than is the class of all those people who do or will or *might* exist. In particular, the class of people protected by the Golden Rule *excludes* the very person the early abortion would cause never to exist at all. That person exerts no moral pull over the agent requiring the agent to comply with the Golden Rule in respect of *that* person. That person does not matter morally. And hence the Golden Rule cannot ground the claim on behalf of *that* person that it would be wrong for the moral agent to fail to bring that person into existence and thereby fail to create additional wellbeing for that person by way of the early abortion.

However, for this objection against Hare to have any effect, things cannot be left so vague. We must say why the possible future person, the person the early abortion would prevent ever existing at all, is to be excluded from the class of those who matters. If that person

is “out,” one what ground? And who is “in?” Who is “on the list?” And who is not? It is when we see that there are no good answers to that question that we also see how strong Hare’s argument really is—and find ourselves compelled at the same time to accept MSMPP.

3. *Who Matters Morally?*

A. *Moral Actualism (alpha)*. Perhaps the reason that the person the early abortion causes never to exist at all does not matter morally is that that person never *actually* exists in the case where the early abortion is *actually* performed, and that the people who matter morally are just the people who do or will exist at the uniquely *actual* world.

But the idea that only actual people matter morally (*Moral Actualism (alpha)*) is problematic. The view that all and only actual people matter morally means that the *permissibility* of an act can depend on its *performance*, contrary to the Rabinowicz principle. Suppose that our theory of permissibility generally emphasizes the importance of creating more wellbeing for people rather than less. What Moral Actualism (alpha) adds to that is just that creating more wellbeing for *actual* people is what is important. We assume Comparability. And we assume as well that we are in a case in which the existence at issue is unambiguously worth having and—to keep things simple—that the early abortion makes things neither better nor worse for anyone *other than* the possible future person the early abortion would cause never to exist at all. Then, Moral Actualism (alpha) in combination with our permissibility theory would seem to instruct as follows: if the early abortion *isn’t* performed at the actual world and the person accordingly at least eventually exists at that world, then the early abortion is wrong, but if the early abortion *is* performed at the actual world so that the person never exists at that world, then it’s perfectly permissible. These results are not inconsistent since we are talking here about two different cases involving two distinct actual worlds. The problem is plausibility:

*performance* does not seem to be the sort of thing that can make an otherwise wrong act permissible or an otherwise permissible act wrong.

Thus, as Caspar Hare points out, the “genocidal adventures of nonactual dictators” are surely wrong whether or not the dictators, their adventures *and* their victims ever exist at all at the uniquely actual world.<sup>14</sup> Moral Actualism (alpha) rejects that point. Hence Moral Actualism (alpha) is false.

B. *Moral Actualism (beta)*. The difficulties with this first answer to the question “who matters morally” suggests a second. Perhaps the reason the person the early abortion causes never to exist at all does not matter morally is that that person never *in fact* exists at the world at which the early abortion is *in fact* performed. Perhaps the people who matter morally are those who do or will exist at whatever world—whether the uniquely actual world or another—the act under scrutiny is itself performed. Moral Actualism (beta) thus takes into account, for the purpose of evaluating Hare’s “genocidal adventures,” the suffering of those people who do or will exist at the world at which those very adventures are undertaken: their suffering *counts against* those adventures wrong, and with little presumably to count in their favor our permissibility theory can easily deem those adventures wrong.

Elizabeth Harman comes very close to adopting this view in the form of what she calls the *Actual Future Principle*. Harman’s discussion makes it clear that she does not use the phrase “actual future” to refer to the people who will exist uniquely *actual* world but rather to express the point that an individual’s moral status fact has to do with whether that individual does or will “actually” exist at a given world.<sup>15</sup> Thus, according to the Actual Future Principle, “[a]n early fetus that will become a person has some moral status. An early fetus that will die while it is still an early fetus has no moral status.”<sup>16</sup> Why? Early fetuses have their moral status only “contingently.”<sup>17</sup> They derive any moral status they have at a world from the moral status,

not of the person they *may or may not* develop into, but rather from the moral status of the person they *will* develop into at that world, if any.<sup>18</sup> If, at a given world, there is, as a result of the early abortion, no such person, then the fetus has no moral status at that world, and the early abortion is permissible. It “requires,” Harman says, “no moral justification whatsoever.”

Harman thinks that the Actual Future Principle supports a “very liberal view about the ethics of abortion.”<sup>19</sup> But the principle itself seems problematic. Suppose that we are in a case in which the woman’s having the early abortion improves the lot of no one at all—not the woman herself, or her already-existing or future offspring, or etc. Then, compare two worlds, one, *w1*, where the early abortion is performed and the fetus is destroyed, and the other, *w2*, where the early abortion is not performed and that same fetus eventually develops into a person who has a life worth living. Harman’s analysis suggests that the early abortion, performed at *w1*, is perfectly permissible. For, according to Harman, since the early abortion is performed at *w1*, there is no future person from whom the early fetus can derive any moral status whatsoever. And it is not wrong, according to Harman, to damage or destroy or injure a thing that has no moral status whatsoever.<sup>20</sup> At the same time, she takes the view that the fetus at *w2* has “at least some moral status,” making it, Harman argues, the kind of thing “we are prohibited from harming” and the kind of thing harm against which requires “moral justification.”<sup>21</sup> Since destroying the fetus in the context of the early abortion *clearly* is to harm it—we are not talking here about a “mere” failure to bring the fetus into existence to begin with—it seems that we are then compelled to say that the choice not to have the early abortion at *w2* is *obligatory*.

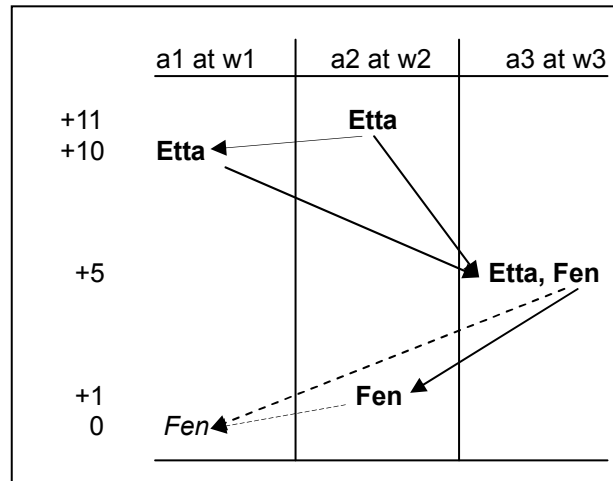
The difficulty is that the result that that choice at *w2* is obligatory has interesting reverberations across the array of the worlds that happen to be accessible to the agent at the time just prior to choice—which array prominently includes *w1*. But if *not* having the early abortion at *w2* is *obligatory*, then *having* the early abortion at *w1* cannot be *permissible* since to say that an act is *obligatory* is to say that no alternative act is itself *permissible*.

So we seem to have an inconsistency. Harman elsewhere clarifies that she does not in fact think that not having the early abortion is obligatory at  $w_2$ . Her position is, instead, that the agent has a *moral reason* at  $w_2$  not to have the early abortion but that the early abortion relative to  $w_2$  is nonetheless *permissible* and requires “no justification” whatsoever. It is perfectly permissible at  $w_2$ , in other words, to destroy this thing that has moral status even though we could have chosen—at no cost to anyone at all—*not* to destroy this thing.

If, however, we are hesitant to disconnect our moral reasons from our moral obligations, we aren’t going to be persuaded that not having the early abortion is not obligatory at  $w_2$ . If we think, that is, that if we have a moral reason *not* to choose  $X$  and no moral reason to *choose*  $X$ —no justification, that is, in favor of the choice of  $X$ —then we are *obligated* not to choose  $X$ , we are not going to be persuaded by that particular response to the charge of inconsistency. We are going to think, instead, that the problem is rooted in the position that one and the same fetus emerging from one and the same history of the world can have some moral status at one possible future and none at all at another.

My concern about the Actual Future Principle can be put in more practical terms. Harman wants to convince her audience that having the early abortion “requires no moral justification whatsoever.” To avoid inconsistency, she must now say that what she means is that the early abortion is permissible at both  $w_1$  and  $w_2$  *but* that the agent has a *moral reason* at  $w_1$  and  $w_2$  in virtue of the fetus’s moral status at  $w_2$  *not* to have the early abortion. Won’t such a position at least leave her audience confused?

C. *A Further Objection Against Moral Actualism.* Moral Actualism, in both its alpha form and beta form, have been severely criticized in recent years.<sup>22</sup> *Addition Plus* represents some of the problems that these views face.



Graph 1: Addition Plus

In this graph, the existential status of a person at a world is indicated by font: bold means that the person *does or will exist* at the indicated world, and italics that the person *never exists* at that world. Downward arrows indicate *losses*—where losses themselves can be incurred *both* when agents bring a person into a lesser existence when they could have brought that same person into a better existence *and* (given Comparability) when agents leave a person out of existence altogether and could have brought that same person into an existence worth having. Whether the arrow is broken or solid also has significance—but that significance we will return to at a later point.

Suppose, then, that a1 is actually performed and that w1 is the uniquely actual world. Etta, and Etta alone, is actual. And Etta is better off in w2 than she is in w1 or w3. Moral Actualism (alpha) implies that Etta, and only Etta, matters morally, and that Etta’s losses, and only Etta’s losses, are to be taken into account for purposes of evaluating each of the alternative acts. That means that Fen and his losses at w2 are to be screened out of the moral evaluation

altogether. Handed these particular moral data regarding losses and their moral significance, our permissibility theory will surely imply that a1 and a3 are wrong and that a2 is obligatory.

But those results seem incorrect. a1 is plausibly permissible—and a2 is clearly wrong. Moreover, even theorists who would evaluate these acts differently—e.g., Totalists—would concede that Fen's plight in w2 clearly at least *bears on* the permissibility of a2 and (still more interestingly) the permissibility of a1 as well. That is: if a1 is permissible, it is permissible because of what a2 does *to* Fen in w2 and notwithstanding what a2 does *for* Etta in w2. Fen's loss at w2 that makes the otherwise wrong a1 perfectly right.

The results we obtain from Moral Actualism (beta) are just as problematic. The beta form at least recognizes that Fen's plight at w2 counts against a2 even if it is a1 that is actually performed. Plausibly, that loss makes a2 wrong. But let's now evaluate a1. For the purpose of evaluating a1, the beta form implies that Fen doesn't matter morally. Fen's plight at w2 has no moral significance for our evaluation of a1; Etta, and only Etta, matters morally for the purpose of evaluating a1. Our permissibility theory can accordingly be expected to instruct that a1 is wrong. As we have seen, however, a1 is very plausibly permissible. Moreover, Fen's loss at w2 surely at least *bears on* the permissibility of a1. Moral Actualism (beta), which says that it doesn't, must accordingly be rejected.

D. *The Prior Existence View*. An entirely different approach to the problem of "who matters morally" is expressed in what Peter Singer calls this the *Prior Existence View*.<sup>23</sup> On that view, when the choice is between bringing a person into an existence of a certain kind and leaving that person out of existence altogether, the only people who matter morally are the people who will exist "prior to the decision we are taking, or at least will exist independently of that decision." Thus the couple who are thinking about having a child should take their own

future pleasure into account, but they need not, according to the Prior Existence View, also “count the likely future pleasure of their children as a significant reason for having children.”<sup>24</sup>


As Singer argues, however, the Prior Existence View can't itself be correct. It goes astray when it completely blinds to the interests of the possible future child. Consider, for example, the case where the couple is thinking of bringing into existence an abjectly miserable child—a child whose life would be *less* than worth living; the child who would be forced to endure the genuinely *wrongful* life. Of course that child's interests *bear on* the couple's decision whether to produce that child: the suffering that child would endure if brought into existence can easily make the couple's otherwise permissible choice to produce that child wrong.

Singer, accordingly, rejects the Prior Existence View. And so should we. Cases involving wrongful life prove that point—and so does Addition Plus; just how is obvious. The difficulty is that Singer then immediately turns to what he calls the “total' view,” according to which “we aim to increase the total amount of pleasure . . . and are indifferent whether this is done by increasing the pleasure of existing beings, or increasing the number of beings who exist.”<sup>25</sup> It's a plus of Totalism that Totalism recognizes the child's interests in never existing at all. But the fact that the Prior Existence view is false is not enough to make Totalism—as Singer no doubt would the first to acknowledge—unless Totalism and the Prior Existence View happen to be our only alternatives. And they aren't.

E. *The Neutrality Intuition and Comparativism.* What John Broome calls the *Neutrality Intuition* takes an approach that is similar to the Prior Existence View, with two important qualifications.<sup>26</sup> First, the Neutrality Intuition is limited to the case where the additional person's wellbeing level falls in the “neutral range.”<sup>27</sup> It thereby avoids the miserable child objection against the Prior Existence View. Second, the Neutrality Intuition by its own terms is introduced as a condition on who matters morally not for the purpose of determining permissibility but

rather for the purpose of determining whether one outcome is morally better than another.<sup>28</sup> Of course, we may also think that the fact that one outcome is *morally better* than another implies that the act that produces the other outcome is *morally wrong* in any case in which agents are in a position to choose between those outcomes (in any case, that is, in which both outcomes are accessible to the agents). If we do think that, then we will consider the Neutrality Intuition to have something to say not just about moral betterness but about moral permissibility as well.

As Broome has clearly shown, however, the Neutrality Intuition cannot be made to work. It implies that the alternatives that differ only in respect of the addition of the person whose wellbeing level is in the “neutral range” are “equally as good as one another.”<sup>29</sup> But Addition Plus suggests that that view is implausible. More plausibly,  $w_2$  is morally worse than  $w_1$ . Broome’s own case is as follows:



Graph 2: Broome’s Argument Against the Neutrality Intuition


The Neutrality Intuition implies that A is morally as good as B is, and that A is also as morally good as C is. Transitivity and Asymmetry then imply that B is morally as good as C is. But we know on other grounds that C is morally better than B is.

Gustaf Arrhenius explores a class of related principles, which he collects under the heading Comparativism. What the principles have in common is the following view:

we should draw a distinction between uniquely and non-uniquely realisable people. The former people exist in only one out of two compared outcomes, whereas the

latter exist in both of the compared outcomes. The idea is that we should take the well-being of non-uniquely realisable people into account in a different way as compared to the well-being of uniquely realisable people.<sup>30</sup>

But Arrhenius shows that the principles that adopt this approach are each problematic in one way or another, with the most straightforward among them, what Arrhenius calls *Strong Comparativism*, failing to take into account the following case<sup>31</sup>:



Graph 3: Arrhenius's Argument Against Strong Comparativism

According to Strong Comparativism, we are to “completely disregard the welfare of *uniquely realisable* people.” But if we exclude those people and their welfare from our considerations, we then seem compelled to say that X is morally better than Y, Y is morally better than Z, and Z is morally better than X.<sup>32</sup> Again, transitivity and symmetry imply that X is morally better than Z. So we have a problem.

F. *Summing Up*. The message throughout this discussion has been consistent. It seems that it will not work to say that some people matter morally and other people don't. The moment we try to simplify the structure of moral law with the aim of subduing an otherwise absurdly stringent procreative obligation—which, given the Thinking Thing Account of Moral Status, in combination with the basic biology of the human embryo and early human fetus, the prohibition against early abortion plausibly becomes—by leaving some person or another “off the list,” we find a scenario in which we are after all compelled to take that very person's needs

and interests into account in order accurately to evaluate *both* the act that imposes a loss on that person at the world where that loss is incurred *and* alternative acts performed at still other worlds. We are compelled to say that all people—you, me and the merely possible—matter morally. We are compelled to accept MSMPP.

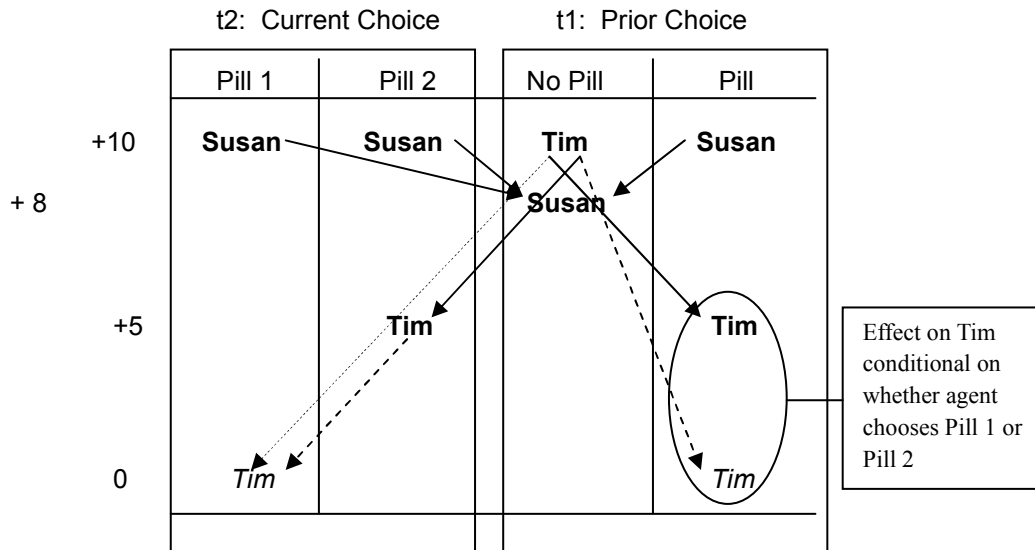
#### 6. *McMahan's Paradox*

We have examined the MSMPP Argument Against Early Abortion and Hare's related argument. Still another argument against early abortion can be found in McMahan's work on the paradoxes of abortion and prenatal injury. That argument comes to the same strong conclusion that Hare's argument does but is grounded in a handful of principles that are far more plausible than Hare's. It's not McMahan's intention to produce a powerful argument against early abortion.<sup>33</sup> I think, however, that that is just what he has done—or at least what his work suggests.<sup>34</sup>

Suppose that we are in a case in which no one's interests, other than those of the fetus itself—and, I would add, the possible future person—are at stake.<sup>35</sup> The agent, a pregnant woman. Susan, faces, at the present moment, two alternatives: she may take Pill 1 and she may take Pill 2. As background, we are to note that Susan has at some prior moment, already chosen to take one pill or the other; she has already, that is, chosen Pill against No Pill.<sup>36</sup> Pill 1 destroys—kills—the fetus. It thus also insures that the child, Tim, the fetus would have developed into never exists at all. Pill 2 injures but does not kill the early fetus. Pill 2, moreover, has a further effect: it causes the future child, Tim, to suffer moderate pain throughout his entire life. Temporally speaking, the No Pill choice is at this stage *behind* the woman; she's already chosen to take one pill or the other; she's already chosen Pill. But the effect of No Pill is nonetheless—as we shall see—relevant to the evaluation of the choices the

agent now has, and we note it here: No Pill neither kills nor injures the fetus—the fetus *and* the later child will *both* exist *and* are perfectly healthy.

Summing up:



Graph 4: McMahan's Case

It seems clear that Pill 2 is impermissible. Even given that Susan has an interest in taking one pill or the other to relieve her own mild chronic pain, we understand that it is wrong for Susan to cause “her own child to suffer [moderate] chronic pain.”<sup>37</sup> Now, in my view, and I believe McMahan’s, what clarifies that Pill 2 is wrong is that Susan at t1 had the choice between No Pill and Pill: that is, the outcome at which Tim both exists *and* is very well off is an outcome that is accessible to the agent at t1.<sup>38</sup> She can’t—let’s assume<sup>39</sup>—go back and change or somehow preempt what she did at t1. But the very fact that she could at t1 have made things so much better for Tim at a relatively small cost to herself, by simply choosing No Pill rather than Pill, reveals just how morally defective the Pill 2 choice in fact is. It shows that Pill at t1, if (at t2) implemented by Pill 2, is itself wrong, in light of the availability of No Pill at t1. (This is not, or at least not *yet*, to say that No Pill is obligatory or even permissible. Instead, it is to say that the

permissibility of Pill at  $t_1$ —assuming, for the moment, that that choice was permissible—depends on how the agent carries out that choice. That is: *if* Pill is permissible under any condition at all, then its permissibility holds only on the condition that the agent’s vehicle for implementing that choice is Pill 1 and not Pill 2. In other words, the agent’s choice of Pill 2 is wrong at any world where the agent makes that choice and hence so—at *that* world—was the agent’s choice of Pill.)

Pill 2 is thus wrong. The interesting point that McMahan then makes is that Pill 1 is what he calls “Pareto inferior” to Pill 2.<sup>40</sup> Pill 2 makes no one, including Susan, any worse off, and Pill 2 is actually *better for the fetus* than Pill 1 is, since it is *better for the fetus* to incur the prenatal injury than it is to be killed. But since we know that Pill 2 is wrong, and Pill 1 is, we’ve now said, “Pareto inferior” to Pill 2, we should infer that Pill 1 is wrong as well.<sup>41</sup>

To understand McMahan’s point, it is important to note that McMahan himself, for purposes of his discussion, assigns some very small amount of moral status to the fetus—*enough* so that the fact that Pill 2 is better for the fetus than Pill 1 has at least some moral significance in its own right.<sup>42</sup> That is why Pill 1 is not *equally as good as* Pill 2 is (to borrow Broome’s phrase) but rather, in McMahan’s view, *inferior to* Pill 2.

But let’s back up. Suppose that we clarify that the fetus in our case is the *early* fetus—the fetus that has not yet developed into a thinking thing, a person. And suppose then that we take the position that that early fetus has *no moral status at all*. Modifying the case in that particular way, we would then, it seems, say that McMahan’s Paretian approach would imply that Pill 1 is equally as good as Pill 2 is. There is one other point: we might accept Comparability—I do; this is a position that McMahan would contest<sup>43</sup>—and accordingly think that Pill 2 is *better for* the child Tim than Pill 1 is. And we might also accept—and I think we must—MSMPP. We may then feel compelled to reinstate McMahan’s original result: that Pill 1 is “Pareto-inferior” to Pill 2, that is, that Pill 1 is *morally worse than* Pill 2, since worse for a person,

Tim, who himself matters morally just as much as you and I do, and better for no one. But whether one accepts this latter point or not, the Paretian point that we are left with is just this: *either* Pill 1 is equally as good as Pill 2 is *or* Pill 2 is Pareto-inferior to Pill 1. We then obtain the very same results that McMahan himself obtains: if Pill 2 is wrong—and surely it is—then, since Pill 1 is *either* equally as good as *or* Pareto-inferior to Pill 2, surely Pill 1 is wrong as well.

In the case at hand, in other words, the early abortion is wrong.<sup>44</sup> Now, one might think the case McMahan describes is highly artificial and hence highly unusual; one might think that the early abortion's being wrong in this unusual case tells us *nothing* about the permissibility of early abortion in any ordinary case.

But that's not so. For in almost any case in which the woman has the option of the early abortion she also has the option of injuring the fetus and impairing the future child. The woman could decline the early abortion but consume a *great* deal of alcohol while pregnant, or take dangerous prescription drugs, or knowingly infect herself with the measles or some other virus that leads to birth defects. She could experiment with Thalidomide early in pregnancy. She could take drugs that trigger labor, for that matter, long before the 40<sup>th</sup> week of the pregnancy, thereby producing a tiny infant struggling for oxygen and nutrition. The list of ways in which a pregnant woman can harm the fetus—and, in doing so, the future child—goes on and on.

That means that the logic we have constructed here, if correct, will apply across the board—the early abortion will ordinarily be wrong, since ordinarily *that third option will be available to any agent who has the option whether to continue the healthy pregnancy or to have the early abortion*: the agent will have the option of impairing or injuring the early fetus and the future child without killing the fetus and without preventing the child from ever existing at all. Provided that the injury is not so severe that the child's existence is not worth having, the availability of the injuring-without-killing option will be enough to show—according to the logic we have constructed here—that killing the early fetus is wrong.

I think, however, that that logic is *not* correct. I want to try now to explain just how it goes astray—and, more generally, how the position that the early abortion is ordinarily perfectly permissible, whether performed for a good reason, a bad reason or no reason at all, can be reconciled against MSMPP.

## 7. *Variabilism*

We tried to answer the question “who matters morally,” and we tried to answer it in a way that would explain the permissibility of early abortion (and indeed nonconception). But that was an unfortunate strategy. We just aren’t going to be able to divide people up according to who matters morally and who does not.

But it doesn’t follow that we can’t divide up their *losses* according to which *losses* matter morally and which do not.

One might have thought that, if *some* of a person’s losses matter morally, then *that person* matters morally; and if *that person* matters morally, then *all of that person’s losses* matter morally. That is something one might have thought, but one wouldn’t really have had any reason to think that. The latter inference, in particular, is suspect. It’s perfectly possible, in other words, that a person can matter morally, but *variably*, in the sense that *some* of that person’s losses matter morally and *some* do not.

Once we begin to think along the lines of dividing up, not *people* according to whether they matter morally or not, but rather their *losses*, then we quickly see that—as the cases themselves suggest—whether a loss has moral significance or not isn’t random; there is a pattern.

Just think: taking Pill 2 imposes a loss on the child Tim, and we agree that it’s a morally significance loss. In other words: the fact that Pill 2 makes things worse for Tim than No Pill does means that Tim incurs a loss under Pill 2, and that loss counts against the agent’s taking

Pill 2. But Pill 1 is very different. It is plausible to think that Pill 1 imposes a loss on Tim relative to Pill 2 and of course an even greater loss relative to No Pill: Pill 1 makes things much worse for Tim than No Pill does and quite a bit worse for Tim than Pill 2 does. But intuitively the loss that Tim incurs under Pill 1 *does not count against* Pill 1 or *in favor of* either No Pill or Pill 2; that loss has no moral significance whatsoever.

Similarly: it's a loss to you if your mother's experimenting with Thalidomide while pregnant with you caused you to suffer an impairment. That loss has full moral significance; it *counts against* what she did. But it doesn't follow that the loss you would have incurred had she left you out of existence altogether would have had any moral significance at all.

These two cases alone suggest a pattern—and a pattern that we can then recognize in the other problem cases. The idea is that the moral significance of a given loss is a function, not of *who* incurs the loss, but rather *where* the loss is incurred in relation to the person who incurs it. Thus: each person, existing, future or merely possible matters morally, just as MSMPP suggests, but each person matters only variably, in accordance with the following principle:

*Variabilism.* For any world accessible to agents at the time of choice, a loss incurred at a world where the person who incurs that loss does or will exist has *full moral significance* for purposes of evaluating both the act that imposes that loss and the alternative acts that avoid that loss. *But* a loss incurred by that same person at a world where that person never exists at all has *no moral significance whatsoever*, for purposes of evaluating the act that imposes that loss or for purposes of evaluating the alternative acts that avoid that loss.

This way of looking at things explains—interprets—MSMPP. It gives MSMPP the credit it deserves. But it doesn't overstate things. Each of us—you, me and the merely possible—incurs many, many losses that have full moral significance. It is fair accordingly to say that we each matter morally. But it doesn't follow that *all* of our losses have moral significance. It is perfectly consistent with the view that some of our losses have moral significance, and that we matter morally, to say as well that some of our losses have no moral significance whatsoever.

Once we begin to divide things up, not in terms of *who* matter morally, but rather in terms of *what losses* matter morally, it is clear what our division should look like. The losses that have moral significance are those that are incurred at worlds where their subjects do or will exist; the losses that aren't are those incurred at worlds where their subjects never exist at all. The point can be made in graphic form. As noted earlier, all the arrows in all the graphs that appear herein land in worlds where the indicated subject incurs a loss. We can note now that—according to Variabilism—the losses indicated by the *solid* arrows have full moral significance (they *count against* the acts that impose them and, in a roundabout way, *in favor of* the acts that avoid them), while the losses indicated by the *broken* arrows have no moral significance whatsoever.

Addition Plus (Graph 1) shows how Variabilism distinguishes itself from its competition. According to Variabilism, Fen's loss *at w2* has full moral significance for purposes of evaluating *both a2 and a1*. Fen's plight at *w2* exerts a moral pull on *both* what agents do at *w2 and*, cross-world, what they do at *w1*. In contrast: Moral Actualism (beta) sees Fen's loss as having significance only for the purpose of evaluating *a2*; while Moral Actualism (alpha) allows that that loss *may* have significance for the purpose of evaluating both *a1* and *a2*, whether that loss in fact has any significance at all will turn on whether *w2* happens to be the uniquely actual world; the Prior Existence View declares Fen's loss irrelevant to both the evaluation of *a1* and the evaluation of *a2*, as does the Neutrality Intuition and Strong Comparativism. At the same time, Variabilism—in contrast to Totalism—also insists that Fen's loss *at w1* has *no moral significance* whatsoever. That loss, according to Variabilism, does not count against *a1*, nor does it provide a basis for the claim that *a2* is in morally better than *a1* is. Of course, *a2* is better for Fen than *a1* is in a *certain* respect: factually speaking, Fen has *more wellbeing* under *a2* than Fen has under *a1*. But *a1* is not better for Fen in any *morally significant* respect since the loss that is

avoided on behalf of Fen under a2 (we can say: the gain that Fen accrues under a2) is itself devoid of moral significance.<sup>45</sup>

Taking, then, all the morally *significant* data into account (Fen's loss at w2, Etta's losses at w1 and w3), and setting all morally *insignificant* data aside (Fen's loss at w1), we should agree, at least, that it is not at all clear that a1 is wrong or that a2 is morally better than a1 is. More plausibly, our permissibility theory will instruct as follows: a1 is permissible notwithstanding the shallow morally significant loss it imposes on Etta, a2 is wrong given the very deep morally significant loss it imposes on Fen and a3 is wrong, given the very deep morally significant loss it imposes on Etta.

The analysis of the case that Broome uses to object to the Neutrality Intuition (Graph 2) is similar. In that case, Variabilism implies that q's loss under B has full moral significance while q's loss under A has none at all. p incurs no losses under any choice and hence no morally significant losses. The otherwise plausible permissibility theory will surely instruct—given those moral data—the intuitively correct result: that A and C are both permissible and B is wrong. We may produce the additional child or not, but if we do produce the additional child, we must create more wellbeing for that child rather than less.

Variabilism also analyzes Arrhenius's case (Graph 3) in a plausible way. That is, it, rightly, makes no distinction at all among the three choices X, Y and Z; it provides no ground for saying that one choice is better or worse than another or that any one choice is wrong. Any otherwise plausible permissibility theory should imply, given the precisely analogous non-loss, morally significant loss and morally insignificant loss we see under each of X, Y and Z, that each of X, Y and Z is permissible.

\* \* \*

Let's now turn to abortion. Variabilism easily defuses the argument against early abortion. We are now in a position to accept that the possible person the early abortion would

prevent from ever existing at all matters morally—as MSMPP suggests—but also to claim that moral law puts no pressure at all on us to bring that person into existence. Yes, that person will incur a loss if the early abortion is performed. Yes, that person matters morally. Yes, many losses that that person incurs in one world or another have full moral significance. However, since the loss that that person incurs under the early abortion is incurred at a world where that person never exists at all, *that particular loss* is devoid of any moral significance whatsoever. It counts *not at all* against the choice of early abortion. If the early abortion is wrong, it is wrong in virtue of what it does to the human embryo or early human fetus *but not in virtue of what it does to the person who never exists at all*. Plausibly, if the early abortion is wrong at all, it is wrong because the Thinking Thing Account of Early Abortion is false.

It might seem that Variabilism faces the same inconsistency worry that we see in Harman’s Actual Future Principle. After all, though the underlying analysis is very distinct, the conclusion Harman reaches about why the early abortion is permissible at the world  $w_1$  at which the early abortion is performed is at least reminiscent of Variabilism’s own. Harman would say that the early abortion harms the fetus at  $w_1$  but that it does not harm anything that has any moral status at all at  $w_1$  since the possible future person never comes into existence at  $w_1$ . Variabilism says, instead, that the early abortion imposes a loss on (we could just as easily say that it *harms*) a certain possible future person at  $w_1$ , but that that loss, or harm, is devoid of moral significance since the person who incurs it never exists at  $w_1$ .

From here, however, the two views diverge. Thus, consider the issue whether the early abortion is permissible at  $w_2$ , where the early abortion is not performed and the possible future person eventually does come into existence. According to Harman, the early fetus at  $w_2$  has a moral status that it derives from that “actual future” person. And that moral status, according to Harman, gives rise to a moral *reason*—and arguably a moral *obligation*—not to have the early abortion since the early abortion destroys something that has moral status. In contrast,

Variabilism avoids any such implication. The loss incurred as a result of the person's never existing at all has no moral significance at  $w_2$ , despite the fact that that person will in fact eventually exist at  $w_2$ . (Similarly, while you matter morally, and the loss you would have incurred had your mother experimented with Thalidomide while pregnant with you has full moral significance, the loss you would have incurred had your parents never brought you into existence at all has, according to Variabilism, no moral significance whatsoever.) The loss the person incurs at  $w_1$  lacks moral significance, in other words, both at  $w_1$  and at  $w_2$ ; it lacks moral significance *across the board*, at each and every world where the person who incurs that loss does exist, *and* at each and every world at which the person who incurs that loss never exists. That the person eventually exists at  $w_2$  does not somehow imbue that person's  $w_1$  loss with moral significance at  $w_2$  that it does not have at  $w_1$ . Still incurred at a world where the child never exists, the loss still has no moral significance whatsoever.

I end with a note on the argument against early abortion that is suggested by one strand of McMahan's work (Graph 4). The objection from Variabilism focuses on McMahan's Paretian analysis. According to that analysis, since Pill 1 is "Pareto inferior" to Pill 2, being worse for the fetus—or, we said, perhaps simply worse for the possible future person—than Pill 2 is and better for no one, then since Pill 2 is clearly impermissible, so is Pill 1. Variabilism would insist that the loss the child, Tim, incurs under Pill 2 has *full moral significance*. After all, there is clearly a loss; the comparison between Pill 2 and No Pill proves that much. And Variabilism implies that that loss has full moral significance since Tim will at least eventually exist at the world where the agent takes Pill 2. At the same time, Variabilism insists that the loss Tim incurs under Pill 1 has no moral significance whatsoever. It's true that there is a loss—Pill 1 is worse for Tim than Pill 2 is. But Variabilism implies that that loss has no moral significance—it cannot, that is, be counted against Pill 1—since Tim never exists at the world where the agent takes Pill 1.

Variabilism thus allows us to construct a very plausible barrier against the view that Pill 1 is “Pareto inferior” to Pill 2 and, accordingly, against the implication that Pill 1 is itself wrong. It provides us, in other words, with a basis for insisting that the bare fact that Pill 2 is, *in respect of wellbeing*, better for each person and worse for no person than Pill 1 is does not on its own establish that Pill 1 is “Pareto inferior,” or morally worse than, Pill 2 is, or that Pill 1 is wrong. The analysis is as follows. The respect in which Pill 2 is better for Tim than Pill 1 is—that is, that Tim has *more wellbeing* under Pill 2 than Tim has under Pill 1—is, according to Variabilism, without moral significance. This is just another way of saying that the loss Tim incurs under Pill 1 as compared against Pill 2 is itself devoid of moral significance. At the same time, the respect in which Pill 2 is worse than No Pill for Tim has full moral significance: Tim’s loss under Pill 2 (compared, of course, not against Pill 1 but rather against No Pill) counts, definitively, against Pill 2. In short, Tim incurs no morally significant loss under Pill 1 and a very deep morally significant loss under Pill 2. In the respect that is itself morally significant—and *not* the merely factual has-more-wellbeing respect—Pill 2 is worse for Tim than Pill 1 is. From that point, it is hard to see how any plausible theory of outcome betterness—or act permissibility—could then produce the result that Pill 1 is Pareto inferior to Pill 2, or that Pill 2 shows that Pill 1 is wrong. The usual Pareto formulation—if one act makes things better for some and worse for none than an alternate act does, then that alternate act is wrong, where *betterness-for* is understood by reference to the concept of *having additional wellbeing* and nothing more—is perfectly fine in the context of “same people” cases. But, in context of “different people” cases—including McMahan “different number” case—its inherent plausibility is lost in the absence of careful interpretation.

8. *Conclusion*

The early abortion prevents a possible person from ever coming into an existence that is, in many cases, worth having. It imposes a *loss* on that person. But the loss itself is imposed at a world where that person never exists at all. For that reason it is, according to Variabilism, devoid of moral significance. It does not bear on the moral evaluation of the early abortion; it does not *count against* that choice; it cannot make the otherwise permissible early abortion wrong.

On the other hand, it is at least plausible that the late abortion ends the life of an existing person—that it imposes a loss on a person at a world where that person, that thinking thing, exists. If that is correct, then the loss that the late abortion imposes will have full moral significance, according to Variabilism; it counts against the late abortion; it is perfectly capable of making the otherwise permissible late abortion wrong.

Variabilism suggests a middle ground on the moral significance of the merely possible: they do matter morally, but only (as do we) variably. The middle ground that Variabilism suggests on the moral significance of the merely possible opens the door to a middle ground on abortion. Variabilism supports the argument that the early abortion is ordinarily permissible, whether it is chosen for a good reason, a bad reason or no reason at all. In contrast, the late abortion is subject to a different analysis—at least given the assumption that being a person and being a thinking thing and mattering morally are all interconnected in the manner that the Thinking Thing Account of Moral Status imagines, and given the further assumption that the late fetus is indeed a thinking thing, a person, so that the abortion is a matter not of leaving a merely possible person out of existence altogether but rather of ending the life of an existing person.

\* \* \*

The main question here has been this. If a merely possible person matters morally, and if early abortion creates less wellbeing for that person when the agents could have created

more, then doesn't the fact that that avoidably lesser wellbeing level is created for that person at least *count against* the early abortion? Can't it, that is, at least in many cases make the otherwise permissible choice of an early abortion *wrong*?

My answer has been that it can't. It's just not the kind of loss that matters morally— notwithstanding the fact that still other losses of equal or greater dimension, or that cut just as deep or even deeper, incurred by that very same person do matter morally.

It's fair to ask why—why one sort of loss matters morally, and the other matters not at all, when both are incurred by people who have full moral significance. My answer to that question is this: dividing up losses in accordance to where they incurred in relation to the person who incurs them explains the cases, while dividing up persons in accordance with their *modal* status—whether they happen to be *actual* or merely possible; or *existing under the act under evaluation* or not; or *existing independently* of how the choice is made—cannot be made to work. Perhaps it's just axiomatic, that each person, you, me and the merely possible, has full moral significance, but that all that that really means is that some of our losses matter morally while others do not, with the moral significance of a given loss itself being a function of *where* that loss is incurred in relation to the person who incurs it.

It's fair to ask as well whether Variabilism does not make the mistake of assigning a moral significance to *losses* to the exclusion of *gains*. That particular objection misunderstands the sense of "loss" adopted for purposes here. For purposes here, each and every *loss* can be paired with a corresponding *gain*: losses occur when agents create less wellbeing for a person when they could have created more; gains occur when agents create more wellbeing for a person when they could have created less.

But there is a deeper objection. According to Variabilism, the loss incurred at a world where the person does or will exist has full moral significance but the gain accrued at a world where the person does or will exist may or may not have any moral significance whatsoever.

The gain has moral significance just in the case in which that gain reverses a loss that has moral significance, but not otherwise. But isn't that arbitrary? Why not say, instead, that, just as losses have moral significance when incurred at worlds where the subject of that loss does or will exist, so do gains have moral significance when accrued at worlds where the subject of the gain does or will exist?

The best response to this objection is a simple one, which is just that the described view is *not the view the view I am proposing here*. The described view is equivalent to the position that *all losses* have full moral significance. Some theorists, of course, may want to defend that view. But it leaves us with the very same procreative obligation—a procreative obligation that is absurdly stringent, given the Thinking Thing Account of Moral Status as an assumption, together with the empirical fact that the human embryo and early human fetus have not yet had that first thought—that it has been my aim here to contest. To put the point another way: my aim here has been to show that MSMPP can be reconciled against the permissibility (in any ordinary case) of the early abortion. I thus for purposes here have no interest in the view that says just the reverse.<sup>46</sup>

---

<sup>1</sup> Narveson 1976, p. 73. The view I will describe here—Variabilism—like other formulations of what is called the *person-affecting approach* will be vulnerable to objections based on the *nonidentity problem*. I have elsewhere argued that that problem should not be allowed to decide the issue of whether we must abandon the person-affecting approach, or at least the core intuition that what is “bad” must be “bad for” someone (Parfit 1987, p. 363) in favor of an approach that is impersonal even in part. The ground for my argument is not that, e.g., the depletion choice or risky policy that Parfit describes, or the slave child contract or pleasure pill ingestion that Kavka describes (Kavka 1982), are not wrong but rather that harm—in an intuitive, comparative sense of that term—is more easily identifiable in those cases that Parfit and Kavka themselves suggest. Roberts 2009c and Roberts 2007.

<sup>2</sup> McMahan rejects the idea that bringing a miserable person into existence or leaving a happy person out of existence could make things better or worse *for that person*. “[T]o cause a miserable person to exist cannot be worse for that person, since ‘worse for’ implies a comparison with an alternative that would be better . . . . [But] *there is never anyone* for whom that alternative is better . . . .” McMahan 2009, p. 59 (emphasis added). See Broome 1999, p. 168, and Bykvist 2007, pp. 339-345 (arguing against Comparability). See also Johannsson 2010 and Arrhenius and Rabinowicz 2010 (arguing in favor of a limited form of Comparability).

<sup>3</sup> Similarly, whether we are modal actualists or not, we should agree that we must be able to find a way meaningfully and truthfully say to say that JFK could have had a third child who was an astronaut but could have been a senator. McMichael 1983. We must, that is, find a way of talking “about” people who never in fact exist at all.

<sup>4</sup> McMahan 1981. Heyd, 1992, pp. xi-xii, 16, 21-31 and 111-115. Thus, according to Heyd, “‘be fruitful and multiply’ should be understood as a blessing rather than a commandment.” The act that brings a person into existence “logically” can’t, he thinks, make things better or worse for that person, since there is *no person* (under the alternate choice) to say that things are worse or better *for*. Comparisons between the scenarios in which the anguished child does and doesn’t exist are problematic because “nonexistence . . . is not a state that can be *attributed* to a subject. It is hardly a ‘state’ at all. . . . Had the plaintiff in the Israeli case [involving wrongful life] not been conceived, would it have been a gain for anyone, would *this* boy have been better off? Similarly, is being born an injury or harm in his case?” (pp. 30-31).

<sup>5</sup> A natural extension of this assumption would be to say that two thoughts are thoughts of one and the same person iff they are bound together by a relation of psychological continuity. The theorist who very much dislikes the metaphysics here, including the idea left open by these assumptions that a late-term fetus may for a certain duration occupy exactly the same space as the person that fetus can be correlated with, can feel free to convert this talk about when persons come into existence to talk about when the individual that is at one time “just” a fetus transitions from a non-person stage, a stage devoid of thought or feeling, to a person stage.

<sup>6</sup> Hare 1975, p. 208.

<sup>7</sup> Hare 1975, p. 208.

<sup>8</sup> Hare 1976, p. 221.

<sup>9</sup> Here I refer to the Principle of Normative Invariance.

<sup>10</sup> Hare 1975, p. 219.

<sup>11</sup> Hare 1975, p. 220.

<sup>12</sup> We should note in passing that such limits are fairly weak—the early abortion, or even the nonconception, is permissible *only if* it maximizes wellbeing on an aggregate basis, something it will only rarely accomplish in the real world, where the child will be born into a situation in which its parents will make sure to create for it a reasonably good life.

<sup>13</sup> Hare 1975, pp. 218 and 221-222.

Totalism is often understood to include two points. The first is that one choice is *morally better than another* just in case it produces a greater amount of *wellbeing* on an aggregate, or *total*, basis than the other does. And the second is that a given choice is *permissible* just in case no alternate choice produces a still greater amount of wellbeing on a total basis. Certainly, Totalism embraces MSMPP (even if it would point the point in very different terms). But theorists who endorse only the first of those two points (for example, John Broome; see Broome (2004), 30-39.) would seem to embrace MSMPP as well: they must concede that that first point—functionally, if not explicitly—assigns to each person—existing, future or merely possible—a moral significance thoroughgoing enough to imply that the failure to take into account the interests of a merely possible person within the context of a given possible future is fully capable of making that future morally *worse* than still another future in which that person's interests are taken into account.

It might be objected that Totalists don't need to say anything at all about *merely possible people as individuals*—and hence can't be charged with assigning any moral significance to those people or their losses. It's true that the Totalist need not be concerned with relations of cross-world personal identity. Nonetheless, any standard calculation of when one possible future (or distribution) is morally better than another will be based on the array of *individual* wellbeing levels contained in each of the two futures to be compared. Thus, when Broome writes “ $g^p$ ,” he is referring to a vector consisting of a person  $p$ 's wellbeing level within a given distribution at each time at which that person exists within that distribution, with the value of the distribution itself being a function of “all the individuals' vectors [at that distribution] strung together.” Broome 2004, pp. 25-26. Functionally, then, *each person in each distribution*—regardless of the fact that some such persons will be existing or future relative to one distribution, and merely possible relative to another—have moral significance, according to Totalism. Importing personal identity, we can then draw the inference that, just as the person has moral significance, so does any difference between the wellbeing level that person has under one distribution and the wellbeing level that same person has under another—so does, I will say, any *loss* that person incurs at one future relative to another.

<sup>14</sup> Hare 2007, p. 503.

<sup>15</sup> Harman 1999, p. 321. On a related topic, Harman thus elsewhere writes (citing Harman 1999): Some people “worry that all actual reasons against harming must stem from actual things that actually have moral status. But this is a mistake. Reasons against a course of action need not rely on what the world is actually like; they can rely on what the world would have been like if the action had been performed.” Harman 2004, p. 107.

<sup>16</sup> Harman 1999, p. 311.

<sup>17</sup> Harman 1999, p. 321.

<sup>18</sup> Harman 1999, p. 312.

<sup>19</sup> Harman 1999, p. 310.

<sup>20</sup> Harman 1999, p. 313.

<sup>21</sup> Harman 1999, pp. 312 and 315. She writes: “It is possible to give a good account of how the very liberal view is compatible with prohibitions on harming early fetuses that will become persons” (p. 315). Harman considers this result a plus of the Actual Future Principle. But since destroying and killing the early fetus *is* to harm the fetus, then she seems to be endorsing the view that, at the world where the early abortion is not performed, the early abortion—like any other harm to the fetus—can be “prohibited” or at least requires justification.

<sup>22</sup> Hare 2007.

<sup>23</sup> Singer 1999, p. 103.

<sup>24</sup> Singer 1999, p. 103.

<sup>25</sup> Singer 1999, p. 103.

<sup>26</sup> Broome 2004, pp. 143-149.

<sup>27</sup> Broome, 2004, p. 146.

<sup>28</sup> Broome 2004, p. 145.

<sup>29</sup> Broome 2004, pp. 147-148.

<sup>30</sup> Arrhenius 2009, p. 290.

<sup>31</sup> Here, Arrhenius cites Temkin. Arrhenius 2009, p. 295 and n. 19.

<sup>32</sup> Arrhenius 2009, p. 293.

<sup>33</sup> His intention, rather, is to produce a paradox: a pair of abortion choices we intuitively think should be evaluated in the same way (that is, either both permissible or both wrong) but find ourselves compelled under a certain Paretian analysis (according to McMahan) to evaluate differently (one permissible and the other wrong). McMahan himself proposes that the way to resolve the paradox is to appeal to a difference in the “order” in which the various alternatives that are available to the agent are compared by the agent; we can then see, he thinks, why the one choice is permissible while the other is wrong. McMahan 2006, pp. 640-41 and 650-651. My view is that the Paretian analysis that he uses to establish wrongdoing in the one case is such that, if it’s correct, we would be compelled to accept wrongdoing in both cases, but that in fact it is incorrect.

<sup>34</sup> See note 33 above.

<sup>35</sup> Making that assumption puts us in the part of McMahan’s Paradox of Sequential Choice that involves the two forms rather than just the single form. McMahan 2006, pp. 650-651. Specifically, we are in the part of the two forms scenario that imagines that the woman has already chosen to take steps to relieve the chronic mild pain she will otherwise suffer. Since her interests are, accordingly, already protected, the only individuals who might have interests at stake are the fetus and the possible future person.

<sup>36</sup> See note 34 above.

<sup>37</sup> McMahan 2006, p. 651. And now I should note how I have modified McMahan’s case: the woman, if she fails to take some pill or another, will suffer chronic mild pain. But even the need to address that chronic *mild* pain would not suffice to justify her taking Pill 2, since the future child under the choice of Pill 2 will suffer chronic *moderate* pain (in McMahan’s case).

<sup>38</sup> McMahan goes so far as to write: “It seems that she ought now to write back, revoking the choice she made on the first form. Indeed, if she were particularly clever, she might have anticipated while reviewing the first form that if she were to check the treatment option, she would then have a decisive reason, on receiving the second form, to choose pill 2. Given that knowledge, it seems that her only permissible option on first form was to choose not to be treated.” McMahan 2006, p. 651.

<sup>39</sup> McMahan suggests that she can; see note 38 above. I do not think we need to assume, however, that she can correct her prior choice in order to show, by reference to that choice, that Pill 2 is itself wrong.

<sup>40</sup> McMahan 2006, p. 641.

<sup>41</sup> McMahan 2006, pp. 640-641.

<sup>42</sup> Indeed, it may not even qualify as an “early fetus,” that is, as a fetus that is not yet a thinking thing.

<sup>43</sup> This is a view that McMahan rejects, however, because he rejects Comparability. We have, however, for purposes here assumed Comparability to be true

<sup>44</sup> That is exactly McMahan’s conclusion—though he also produces a line of reasoning to the conclusion that Pill 1 is permissible. It is those two lines of reasoning together that, he argues, create the paradox. See note 33 above.

<sup>45</sup> Even so, it may seem that consistent with Variabilism we are still compelled to declare a1 wrong on the Paretian ground that a2 is “better for all and worse for none” than a1 is. It’s true that, factually speaking, each of Etta and Fen has more wellbeing under a2 than each has under a1. And it’s true that Etta’s loss (unlike Fen’s) under a1 has full moral significance. So we can say that a2 is better for Etta than a1 is in a morally significance respect. The barrier that Variabilism allows us construct against the deeper conclusion—the conclusion, that is, that a2 is morally better than a1 is and hence that a1 is wrong—is just this: a2 is worse for Fen than a1 is in a morally significant respect, and it is *that* respect (*not* the merely factual respect) that is relevant for any determination of overall moral betterness. More specifically: Variabilism implies that the loss Fen incurs under a2 (and in proving that there *is* such a loss we compare a2 not against a1, of course, but against a3) has full moral significance, whereas the loss Fen incurs under a1 has none at all. In the morally significance respect, in other words, a2 is (given the option of a3) worse for Fen than a1 is.

<sup>46</sup> I presented some of the arguments in this paper—in particular, the arguments against different answers to the question “who matters morally” and the argument in favor of Variabilism—in Roberts 2010. The analyses of both McMahan’s and Harman’s positions that I presented in that earlier work have been expanded, and the discussion of McMahan’s Paretian analysis corrected. For their comments on earlier drafts of this paper and a related paper, I am very grateful to Fred Feldman, Elizabeth Harman, Ingmar Persson, Peter Singer and David Wasserman.

### Bibliography

- Arrhenius, Gustaf. 2009. "Can the Person Affecting Restriction Solve the Problems in Population Ethics?" In *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, eds. Melinda A. Roberts and David T. Wasserman. Springer, pp. 289-314.
- Arrhenius, Gustaf and Wlodek Rabinowicz. 2010. "Better To Be than Not To Be?" Copy on file with author.
- Broome, John. 1999. *Ethics Out Of Economics*. Cambridge University Press.
- 2004. *Weighing Lives*. Oxford University Press.
- Bykvist, Krister. 2007. "The Benefits of Coming into Existence." *Philosophical Studies* 135: 335-362.
- Hare, Caspar. 2007. "Voices from Another World: Must We Respect the Interests of People Who Do Not, and Will Never, Exist?" *Ethics* 117: 498-523.
- Hare, R.M. 1975. "Abortion and the Golden Rule." *Philosophy & Public Affairs* 4(3): 201-222.
- Harman, Elizabeth. 2000. "Creation Ethics: The Moral Status of Early Fetuses and the Ethics of Abortion." *Philosophy & Public Affairs* 28 (4): 310-324.
- . 2004. "Can We Harm and Benefit in Creating?" *Philosophical Perspectives* 18: 89-113.
- Johansson, Jens. 2010. "Being and Betterness." *Utilitas* 22(3): 285-382.
- Kavka, G. 1981. The Paradox of Future Individuals. *Philosophy & Public Affairs* 11: 93-112.
- McMahan, Jeff. 2006. "Paradoxes of Abortion and Prenatal Injury." *Ethics* 116(4): 625-55.
- . 2009. "Asymmetries in the Morality of Causing People to Exist." In *Harming Future Persons*, eds. Melinda A. Roberts and David T. Wasserman. Springer, pp. 49-68.
- McMichael, Alan. 1983. "A Problem for Actualism About Possible Worlds." *Philosophical Review* 92(1): 49-66.
- Narveson, Jan. 1976. "Moral Problems of Population." In *Ethics and Population*, ed. Michael Bayles. Schenkman Publishing Co., pp. 59-80.
- Parfit, Derek. 1987. *Reasons and Persons*. Oxford: Oxford University Press.
- Roberts, Melinda A. 2002. "A New Way of Doing the Best that We Can: Person-Based Consequentialism and the Equality Problem." *Ethics* 112(2): 315-350.
- . 2003. "Can It Ever Have Been Better Never to Have Existed At All? Person-Based Consequentialism and a New Repugnant Conclusion." *Jo. of Applied Philosophy* 20(2): 159-185.
- . 2007. "The Nonidentity Fallacy: Harm, Probability and Another Look at Parfit's Depletion Example." *Utilitas* 9(3): 267-311.

—. 2009a. "What Is the Wrong of Wrongful Disability? From Chance to Choice to Harm to Persons." *Law and Philosophy* 28(1): 1–57.

—. 2009b. *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, ed. with David T. Wasserman. Dordrecht: Springer.

—. 2009c. "The Nonidentity Problem and the Two Envelope Problem: When Is One Act Better for a Person than Another?" In *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, eds. Melinda A. Roberts and David T. Wasserman. Dordrecht: Springer, pp.201–228.

—. 2010. *Abortion and the Moral Significance of Merely Possible Persons: Finding Middle Ground in Hard Cases*. Dordrecht: Springer.

Singer, Peter. 1999. *Practical Ethics*, 2d. ed. Cambridge University Press.