Chapter 1
Introduction

1.1 Goals

1.1.1 I have two main goals in this book. The first is to give an account of the *moral significance of merely possible persons* – persons who, relative to a particular circumstance, or possible future or *world*, *could* but in fact *never do* exist.[1] I call that account *Variabilism*. My second goal is to use Variabilism to begin to address the problem of abortion.

1.1.2 We ought to do the best we can *for people*. And we consider this obligation to extend to people who are, relative to a world, *existing* or *future*. But does it extend to *merely possible people as well*? And, if it does, then does it extend to making things better for them by way of *bringing them into existence*? If we say that surely it *doesn't*, does that then mean that our obligation to do the best we can for people does *not*, after all, extend to the merely possible – that the merely possible do *not* matter morally? But if the merely possible do *not* matter morally, then doesn't that mean that it would be permissible for us to bring them into miserable existences – and even *obligatory* to do just that – in the case where bringing the merely possible into miserable existences creates additional wellbeing for existing

[1]References to *merely possible persons* and, later on, to persons who do exist – *existing persons* – and persons who will exist – *future persons* – can succeed only if they are understood in relation to particular circumstances, or possible futures or *worlds*. Thus, a person may be *merely possible* (or *future* or *existing*) relative to one world but not relative to another. Accordingly, I relativize those terms to worlds (at least implicitly) throughout this work.

Some philosophers would rather not talk about the merely possible as though they were alive and well and sleeping on the couch in the next room (and would that they would leave). The concern is that such talk will predispose us to think that the merely possible matter morally. But talk about the merely possible is not *inherently* question-begging. At least, it does not beg any question of interest here. And making talk about the distinct ways in which the future might unfold and the distinct populations that do or will exist within those futures – within those *worlds* – out of bounds also makes it very cumbersome to say what we *do* need to say about the merely possible. It would be like trying to say *exactly* what it *means* to say *all men are mortal* without availing oneself of the quantifier. See, for example, Russell (1903), pp. 18–23.

M.A. Roberts, *Abortion and the Moral Significance of Merely Possible Persons: Finding Middle Ground in Hard Cases*, Philosophy and Medicine 107,
DOI 10.1007/978-90-481-3792-3_1, © Springer Science+Business Media B.V. 2010

and future people? Doesn't that mean that doing anything *less* than our best for existing and future people is *always wrong* – even when the reason we do less than our best for them is to avoid imposing miserable existences on the merely possible? But if we recognize *that* as a result that we cannot accept, then doesn't that mean that we are compelled, after all, to conclude that our obligation to do the best we can for people extends to the merely possible?

Breaking out of this tail-biting line of questioning requires some work. In this book, I try to do that work by developing a clear and plausible account of the moral significance of merely possible persons. That account – *Variabilism* – rejects the idea that the merely possible *always* matter morally. But it also rejects the idea that they *never* matter morally. Instead, it claims that the merely possible matter morally but also that they matter *variably* – and that that, as it happens, is *exactly* how all of the rest of us matter morally as well.

The cases that we examine in what follows suggest just such a variability. As we shall see, however, strictly speaking that variability extends not across the domain of *persons* but rather across the domain of *diminutions in wellbeing*, or *harms*, that persons, whether existing, future or merely possible, suffer – that is, across the domain of *losses* that they incur. Accordingly, we don't say that some *people* matter morally and others do not. Instead, where *loss* is understood as shorthand for any case in which agents could have created more wellbeing for a given person and instead create less, we say that some *losses* incurred by that person matter morally and others do not. Some losses incurred by a given person have, that is, *moral significance* for purposes of evaluating the acts that impose those losses and their alternatives and others do not.

The variability that the cases suggest is, moreover, not at all random. There is a pattern. The moral significance of any particular loss is a function of just *where* that loss is incurred in relation to the person who incurs it. More specifically: a loss incurred at a world where the person who incurs that loss does or will exist, according to Variabilism, has *full moral significance*, while a loss incurred by that very same person at a world where that person never exists at all has *no moral significance whatsoever*.

1.1.3 My second goal is to begin to address the problem of abortion. It is useful, I suggest, to explore abortion as a problem in population ethics.

For purposes here, I use the term *early abortion* to describe the abortion that leaves a person out of existence altogether. And I use the term *late abortion* to describe the abortion that ends the life of an already-existing person. An early abortion, in other words, imposes a loss on a person who never exists at all at the world where the early abortion is performed. It imposes a loss on a *merely possible person*. In contrast, a late abortion imposes a loss on a person who has already come into existence at the world where the late abortion is performed. It imposes a loss on an *existing* person. What makes the terms *early* and *late* abortion apt is a certain assumption – that *thinking* and *being a person* come together – in combination with certain empirical facts about just when during the process of development the human organism begins to think. What gives the distinction between the early abortion and the late abortion its moral relevance is a second assumption. The second assumption

is that being a *person*, that is, a *thinking thing*, and being the kind of thing in respect of which we have moral obligations (*mattering morally*) themselves come together. What makes this second assumption plausible may simply be that it is in the case of thinking things that it *matters to them* (whether they know it or not) that we treat them in one way rather than another.

Certainly, these assumptions raise questions that warrant books of their own. For purposes here, however, I am interested in another set of questions. First, I want to know whether the loss incurred by the merely possible person when that person is left out of existence altogether *counts against* the choice of early abortion. Can, in other words, that loss make the *otherwise permissible early abortion wrong*? And, if we say it can't, I want to know, second, whether we can take that position without committing ourselves to the (highly problematic) further position that the merely possible do not matter morally *at all*.

Citing Variabilism, my argument will be that the early abortion, in any ordinary case, is perfectly permissible if it is what the pregnant woman herself wants, and that it can be determined to be permissible under a *very* summary moral analysis. So, similarly, can we determine that the choice not to conceive a child, or to conceive a child in vitro but not arrange for the newly created embryo to develop *in utero* into a person, is permissible in any ordinary case. For, according to Variabilism, the effect the early abortion has on the person it relegates to the class of the merely possible – the loss of wellbeing the merely possible person incurs when that person is never brought into existence to begin with – has *no moral significance whatsoever* for purposes of evaluating the early abortion or its alternatives. We can then simply observe that any otherwise plausible permissibility theory will imply, in any ordinary case, that the early abortion is permissible.

As we shall see, other accounts as well suggest very similar results for this very same class of cases. But those accounts take a radically *exclusive* approach in connection with the merely possible. They make the mistake of asserting that the merely possible person does not matter morally *at all*. And that mistake, in turn, leads the exclusive approach to generate a slew of results that are highly counterintuitive or conceptually problematic or both. Variabilism, in contrast, by acknowledging that the merely possible *do* matter morally – if variably – avoids those results. Yet Variabilism avoids those results without adopting a radically *inclusive* approach in respect of the merely possible. Variabilism thus retains the ability to explain just why the loss incurred by the merely possible person when we choose to leave that person out of existence altogether cannot count against that choice.

I will also be interested here in the problem of late abortion. Variabilism, I will argue, lends support to the view that, ordinarily, late abortion is wrong. The loss that is imposed by the late abortion, according to Variabilism, has *full moral significance*. It's a loss, after all, that is incurred by an *existing* person. And it's often a very great loss – the loss of the person's entire future life. Of course, there are morally significant losses on the other side as well – the pregnant woman and others may well incur losses if the pregnancy is continued. In such cases, the late abortion constitutes a *tradeoff scenario*. Those are always a little complicated, and the analysis cannot be quite as summary as it is in the case of early abortion. Still, in the case

of late abortion, the tradeoff is often relatively simple. Under one choice, one person faces a very great morally significant loss. Under the other choice, other persons face losses that have full moral significance but are not, in any ordinary case, nearly so great. The upshot is that, in any ordinary case, the late abortion is going to be very hard to justify. Once the extent of the loss and its moral significance are both recognized, the otherwise plausible permissibility theory can be expected to declare the late abortion wrong in many cases.

As noted, these results are subject to the various assumptions about what counts as a person. The first assumption is that being a *person* and being a *thinking thing* are connected. I take that assumption to mean that, toward the beginning of the pregnancy, the human embryo or fetus, because it is not a thinking thing, is not a person. But late in the pregnancy, the human fetus, because it is a thinking thing, is a person. At that stage, moreover, it seems plausible that the fetus has acquired enough of a *capacity to survive* from one moment to another that the late abortion's *not* being performed would make things better for that person and not just remove someone from existence whose going out of existence was imminent in any case. The second assumption is that being a person and mattering morally – if variably – come together.[2]

These assumptions, together with Variabilism and our independent sense of how the morally significant loss *bears on* the question of a particular act's permissibility, make the question of *when* the developing human (or non-human) organism begins to think in a way that is critical to that thing's mattering morally very important for purposes of evaluating any given abortion choice. According to the rough timeline I will later describe, the *thinking* that signals the coming into existence of a person, along with the *capacity to survive* that means that the abortion's *not* being performed would make things better for that person, does not take place until at least half-way through the pregnancy.[3] Moreover, the period of time between the moment at which we have *thinking plus the capacity to survive*, and the moment of *birth*, may well itself

[2] We can leave open whether that obligation itself is rooted in an obligation to make things better for certain things *for their own sake* or an obligation to make things better for certain things *for the sake of the universe*. See part 1.3.3 below.

[3] Of course, much depends on what we mean here by *thinking*. But it seems clear that mere electrical activity is not enough. Nor are mechanical responses to pain stimuli. See Peter Singer (1994), pp. 104–105. At the same time, for purposes here we want a concept of *thinking* that is broad enough to extend to conscious beings whose thoughts are quite unsophisticated. They need not have a life plan or a clear concept of themselves as continuing beings. Moreover, we should recognize that thinking often takes place beyond our own conscious awareness. I thus take it for granted that the cognitive activity and degree of continuity we typically see in both the human neonate and in many non-human animals is enough to show that the destruction of those individuals is a matter of ending the life of an already-existing *person* and not the choice *not* to bring a new person into existence to begin with. More generally, the guide for what counts as *thinking* for purposes here is the idea that the kinds of things that matter morally – that are *persons* – are the kinds of things such that it *matters to them* (whether they think so or not) whether we treat them in one way rather than another. If there's enough cognitive activity that that latter is the case, then that activity should be considered thinking, for purposes here. The connections among being a person, mattering morally and thinking are discussed in more detail in parts 1.6.5, 5.2 and 5.3 below.

be substantial – perhaps weeks in duration if the pregnancy is full-term. The timeline I provide, however, is rough – and it's really only a good guess. For my aim in this book is to get the moral significance of the *merely possible person* as it relates to abortion exactly right. It is not getting the timeline itself exactly right.

1.1.4 The matters I have dealt with by assumption here are examined in many old debates on abortion in considerable depth. I think the assumptions I make here are highly plausible. However, even if they are accepted as fully *correct*, on their own they will not even begin to resolve the newer – or at least off-grid – debate having to do with abortion that I want to address here.

That newer debate is triggered by what is really just an argument *against* early abortion – and non-conception and conception followed by non-fruition. The argument is this. Suppose that we become convinced by the cases that it is false that merely possible persons don't matter morally and neither do any of their losses. Suppose that we become convinced that it cannot really be wrong for us to impose some slight loss on existing or future people in a case where the only way we can make things better for them is to bring one or more merely possible persons into existence and treat each such person badly. We will then seem compelled to say as well that the losses that merely possible persons incur are, after all, morally significant. We seem compelled to say that those losses *bear on* the permissibility of the acts that impose those losses and the alternatives to those acts that avoid those losses. And we seem in turn compelled to say as well – by simple universal instantiation – that the loss the merely possible person incurs when that person is never brought into existence to begin with is morally significant and must be *counted against* the early abortion.

Now, it's true that this position does not absolutely preclude the result, in a given case, that the early abortion is permissible. But the analysis will be arduous and, it seems, the result often will be that the early abortion is wrong. After all, the loss that *counts against* the early abortion will be so much graver than anything on the other side that might count *in its favor*.

Variabilism challenges this argument. According to Variabilism, some of the losses incurred by the merely possible have no moral significance at all while still others have full moral significance. We can say that the loss the merely possible person would incur, were we to bring that person into existence and treat that person badly, has full moral significance and counts against that choice. For that loss would be incurred at a world where that same person does or will exist. And we can say at the same time that the loss that same person incurs as a result of the early abortion has no moral significance whatsoever. For that loss is incurred at a world where that person never exists at all.

A variable treatment thus justifies the summary analysis in favor of early abortion – and non-conception, and conception followed by non-fruition. But it does so without putting any pressure on us at all to declare ourselves *obligated* to bring the merely possible into existence and then proceed to treat them in outlandishly bad ways.

1.1.5 The question of the moral status of merely possible persons and the question of the permissibility of early abortion are thus closely connected. I am ordinarily not a big fan of the middle ground. But Variabilism undeniably occupies a middle ground

between two very extreme positions, one that *radically includes* and another that *radically excludes* the merely possible among those who matter morally. According to Variabilism, each person – existing, future or merely possible – incurs *some* losses that have moral significance and *others* that don't, with the moral significance of any loss being a function of *where* that loss is incurred in relation to the person who incurs it. More specifically, a loss incurred by a person at a world is morally significant for purposes of evaluating the act that imposes that loss and each of that act's alternatives *if* that person does or will exist at that world. But the loss incurred by any person – including you and me – at any world where that person never exists at all is completely devoid of any moral significance whatsoever.

This middle ground on the moral significance of merely possible persons opens the door to a middle ground on abortion. Those who think abortion is morally permissible throughout the pregnancy are surely most concerned to have that fact recognized in connection with *early abortion*. And those who disdain abortion at any stage of pregnancy for any reason are surely most concerned to see abortion recognized as impermissible late in pregnancy. My proposal reflects both those concerns. It, in effect, asks each of the two camps to settle for the ground that is most important to it and concede the ground that is most important to the opposition.

1.2 Organization of Book

1.2.1 Do *merely possible* persons matter morally? In evaluating any given act for its permissibility, is a *part* of what is important how the merely possible are affected by that act and its alternatives? Or, alternatively, are persons who *do* or *will* exist the only persons who matter morally? Is it just their needs and interests we must take into account in calculating what we ought to do?

Why these are interesting and important questions, in their own right and in connection with abortion, is outlined in this Chapter 1. My suggestion will be that we can think most productively about these questions if we think not in terms of *who* matters morally but rather in terms of *what losses* matter morally. We can, accordingly, rephrase the critical questions. Do the *losses* the merely possible incur when we leave them out of an existence worth having *count against* the acts that impose those losses? And do those same losses, in a roundabout way, *count in favor of* the alternative acts that *avoid* them?

1.2.2. In Chapter 2, I examine a handful of competing rules on just when the losses incurred by any merely possible person have moral significance and when they do not. Now, those rules do not, on their own, generate any *permissibility* results at all. Rather, each rule simply generates a collection of *moral data* consisting of certain facts that, according to the rule, *bear on* permissibility. Each rule, in other words, does no more than offer a competing picture of which losses are to be taken into account in determining what we ought to do and which losses are *not* to be taken into account in determining what we ought to do.

But, for purposes of testing, it is permissibility that we often have some clear ideas about. And, at least eventually, it is permissibility results that we will need for purposes of giving an account of abortion. Accordingly, in order to test the rules and apply the rule we like to the case of abortion, we will need permissibility results. We will need to see how a given rule on when losses have moral significance (a *loss rule*) combines with principles (*permissibility principles*) that instruct, on the basis of data generated by that rule, which acts are permissible and which are not.

Happily, we can articulate a handful of reasonably plausible permissibility principles that we can accept *regardless* of our views on just which loss rule is correct. These principles together make up what I will call the *otherwise plausible permissibility theory*: the theory that, but for whatever loss rule we happen at the moment to have combined that theory with for purposes of testing, is otherwise plausible.

The first order of business in Chapter 2 is thus to describe that theory, which is also summed up in Appendix A. I shall also need, starting out, to address certain other preliminary matters. For example, before we can say whether a given loss is morally significant or not, we need to say something about loss. For purposes here, I accept a *maximizing* account of loss – an account that finds *loss* whenever agents (by act or omission) create less wellbeing for a person when they could have created more wellbeing for that same person.[4] A bit more needs to be said,

[4] The maximizing account of loss I adopt for purposes here is described in parts 2.2.1–2.2.3 below. Briefly, my account provides that a person – existing, future or merely possible – incurs a *loss* whenever agents (by act or omission) create less wellbeing for that person when they *could have created more* wellbeing for *that same person*. To say that a person incurs a loss is not to say that that person has been wronged, or that a wrong has been done.

This maximizing, comparative account of loss rejects both a *temporal* approach to loss (a person incurs a loss *only if* that person is worse off at a later than at an earlier time) and a *counterfactual* approach (a person incurs a loss *only if* that person would have been better off had the act under scrutiny not been performed). It also rejects a *non-comparative* approach to loss – an approach that defines loss by reference either to a list of specific bad effects (e.g. substantial pain) or to a threshold (e.g. a minimally decent wellbeing level). Finally, it accepts the concept of netting benefits generated by an act for a person against burdens. On my account, we cannot conclude, from the fact that burdens have been generated by the performance of a particular act, that that act imposes a *loss* on a person at a world. A vaccination can impose pain, but that the vaccination imposes some pain does not mean that the vaccination imposes a loss; pain can be imposed even in a case where wellbeing has been maximized.

I also accept, for purposes here, that one way of causing a person to incur a loss – one way, that is, of creating less wellbeing for a person rather than more – is to leave that person out of an existence worth having. Some theorists take the position that a person *cannot* incur a loss at a world where that person *never exists at all*. It seems to me, however, that the concerns that lead to that position can be overcome, and that comparisons between a person's wellbeing level at a world where that person never exists at all against that same person's wellbeing level at a world where that person does or will exist are fully cogent. For a brief discussion of this issue and references to work in this area, see part 2.2.2 and note 45 below.

Some theorists might think that the *nonidentity problem* means that there will be many cases (A) where we *want* to say – and where naively think we *can* say and *need* to be able to say, if we want to make loss central to an account of wrongdoing – that a loss *bears on* the permissibility of a given act but (B) that in fact involve no harm, or loss, at all. We can see how agents could have brought a *nonidentical* person into a better existence, but we cannot – or so the nonidentity problem asserts – see how agents could have brought that *very same person* into a better existence.

however, especially in light of the fact that, to avoid begging any important questions, we shall need an account of loss that recognizes that leaving someone out of an existence worth having is to impose a *loss* on that person. Only then can we open-mindedly ask the next question: is that loss morally significant?

With the preliminaries out of the way, we then examine, also in Chapter 2, the competing rules on when losses incurred by the merely possible have moral significance and when they do not.

The loss rule that I will call *Inclusion* proposes a *radically inclusive position* regarding which losses are morally significant and which are not. *Classic Utilitarianism* (*Totalism*) is a paradigm example of a view that adopts Inclusion. On that view, agents are obligated to make the choice that maximizes wellbeing on a *total*, or *aggregat*e, basis. Whether they do that by creating additional wellbeing for existing or future people, or by creating additional well-off people, is immaterial. The upshot is that the loss a person incurs when that person is left out of existence altogether has full moral significance.

According to Inclusion, the merely possible matter morally just as you and I do, and all of the losses they incur, just like all of the losses you and I and each other existing and future person incur, have moral significance for purposes of evaluating the acts that impose those losses and the alternative acts that avoid those losses.

I argue that the cases show that Inclusion, in combination with the otherwise plausible permissibility theory, is completely untenable. Inclusion fails to take into account a critical, intuitive moral distinction that can only be described as axiomatic. Inclusion must, accordingly, be rejected.

Still another loss rule, what I will call *Exclusion*, is, when combined with the otherwise plausible permissibility theory, roughly equivalent to the view that is sometimes called *Moral Actualism*. Exclusion proposes a *radically exclusive position* regarding loss. According to Exclusion, *merely possible* persons and *existing*

In this book, I discuss the nonidentity problem only briefly, in part 2.2.5 below and in Appendix B. For purposes here, I mainly set it aside, in part because I have discussed it in detail elsewhere (e.g. Roberts 2009c, 2007) and in part because I believe that – *properly understood* – it undermines *nothing* we want or need to say here. My basic point is that the type of nonidentity problem considered among the most telling relies on a fallacy. Cases where I argue that that is so include (among others) Parfit's depletion and risky policy examples, Kavka's slave child and pleasure pill examples and cases involving historic injustices. Once we recognize the fallacy, I argue, we can and must recognize that losses have been incurred and harms imposed. Moreover, in the rare case where we are dealing with a type of nonidentity problem that validly establishes that the act under scrutiny harms and is bad for *no one* – where, for example, a genetically impaired child is brought into an existence that is unambiguously worth having and the choice to bring that child into existence is *maximizing* for that child *and* for each other existing or future person – we are also dealing with a type of nonidentity problem that never clearly establishes that a wrong has been done.

Thus, it seems to me that we should not be too concerned with the nonidentity problem in constructing our moral theory. It is *not* that I think that the choice of depletion or the risky policy, or the choice to enter into the slave child contract or to take the pleasure pill, is permissible. It's rather that we can discern harm, that is, *loss* (and in a comparative, intuitive sense of those terms) in each of those cases. I thus argue that we can resolve the nonidentity problem *without* making the merely possible matter morally *more* than they ought and *without* disconnecting wrongdoing from acting in ways that are bad for some person or another. The key is to take the various types of nonidentity problem one by one and not try to resolve them all in one blow.

or *future* persons are in very different moral boats. *They* don't matter morally, but *we* (existing and future persons) do. Accordingly, while the losses *we* incur have full moral significance, the losses *they* incur have none at all.

Now, as we shall see, Exclusion itself comes in two forms, what I will call its *Alpha* and *Beta* forms, with the Beta form representing an attempt to address some of the issues created by the Alpha form. (Moral Actualism, similarly, is said to come in two forms, *strong* and *weak*.)

I will argue, however, that Exclusion in both its forms is just as untenable as Inclusion is. Exclusion Alpha, in combination with the otherwise plausible permissibility theory, suggests that performance is a determinant of permissibility – that an act can be *permissible*, or even *obligatory*, if *unperformed* but *wrong if performed*.[5] And Exclusion Beta suggests that an act can be *permissible* when its alternative is *obligatory*.[6] Exclusion Alpha seems conceptually unstable while Exclusion Beta on the face of things seems to be flatly inconsistent.

The very gravity of the difficulties that we see in Exclusion may well make us wonder whether we were right to reject Inclusion to begin with.

But we were, for there is a *middle ground* between Inclusion and Exclusion – and thus no need to think that rejecting Inclusion means we are bound to accept Exclusion. According to that middle ground – *Variabilism* – we all – existing, future *and* merely possible – *matter morally*, but we all *matter variably*.

According to Variabilism, each person, whether existing, future or merely possible, incurs *some* losses that have moral significance and *others* that don't, with the moral significance of any loss being just a matter of *where* that loss is incurred in relation to the person who incurs it. That is:

Variabilism:

> Any loss incurred in any circumstance, or possible future or *world*, by *any* person has moral significance for purposes of evaluating the act that imposes that loss or any alternative to that act *if and only if* the person who incurs that loss does or will exist *at the world at which that loss is incurred*.

Variabilism has some important features. The conceptual difficulties we see in Exclusion are rooted in its attempt to say that some people do not matter morally, and neither do any of their losses. Variabilism avoids those conceptual difficulties by taking the view that all persons, existing, future and merely possible, matter morally, and in exactly the same ways.

Inclusion has that last feature in common with Variabilism. But Inclusion then blinds itself to the distinction between (A) the loss the *merely possible* child incurs when we choose not to bring that child into existence to begin with and (B) the loss the real, live, flesh and blood, *existing* child incurs when we create less wellbeing for that child when we could have created more. Variabilism does not do that.

---

[5] Exclusion Alpha, in other words, violates Wlodek Rabinowicz's Principle of Normative Invariance. See note 61 below.

[6] This result violates a standard deontic axiom that asserts that if one act is obligatory, its alternatives cannot be permissible. See part 2.2.4 below.

According to Variabilism, while we all matter morally, we all matter variably. Thus, the merely possible child matters morally, but the loss that child incurs when he or she is left out of existence altogether has *no bearing* on the permissibility of what we have done and, accordingly, can't make the otherwise permissible act wrong. It's the loss in relation to where it is incurred, not the child, that makes that so. It's the loss being incurred at a world where the child never exists that makes it devoid of moral significance. The existing child (of course!) matters morally as well. And the loss that child incurs, when we create less wellbeing rather than more for that child at a world where that child exists, has full moral significance and can easily make the otherwise permissible act that imposes that loss wrong. But again it's the loss and where it is incurred, not the child, that makes that so.

Now, Variabilism does not exhaust the middle ground between Inclusion and Exclusion. A position that John Broome calls the *Neutrality Intuition* and a related position that Peter Singer calls the *Prior Existence View* both can be construed as occupying middle ground as well. The test for determining whether one choice is just as good as another that is proposed by the Neutrality Intuition sets aside the interests of any person who does or will exist under the one choice but not the other – the interests, that is, of any additional person whose coming into existence is at stake – provided that person's existence falls within a "neutral range." Similarly, the Prior Existence View asserts that the determination of the permissibility of any particular choice is to be made by reference *only* to how well off the persons are who do or will exist *independently* of that choice – thus again setting aside the interests of the additional person.

But Broome and Singer argue – correctly, I think – that the Neutrality Intuition and the Prior Existence View are unacceptable. What I want to point out for purposes here is just that Variabilism is a quite different view. We can, that is, easily reject the Neutrality Intuition and the Prior Existence View, as Broome and Singer argue that we must. But there will be nothing in those arguments that suggests that Variabilism is itself defective. Just the reverse: Variabilism helps us articulate how those views go awry.

The notion that the interests of the additional person are to be set aside for purposes of determining whether that person ought to be brought into existence or not is – like Exclusion – sometimes associated with what is called the "person-affecting," or "person-based," intuition. According to the person-based intuition, an act that is wrong must make things worse for, or harm, or impose a loss on, some existing or future person. The association, however, between the person-based intuition and Exclusion is an unfortunate one. Thus, the Neutrality Intuition regards the addition of a person to a world as very often "*morally neutral*."[7] So does the Prior Existence View. It is, as we shall see, that very fact that in the end shall mean that both views must be rejected. In contrast, Variabilism regards the addition of a person to a world

[7] Broome (2004), p. 143 (emphasis added). I here compress (in Broome's terms) the *deontic* reading and the *teleological* reading of the person-based intuition. One pertains to the evaluation of *acts*; the other to the two-place betterness relation that we may think obtains between *outcomes*, or *worlds*, or *acts* performed at worlds. See Broome (2004), pp. 140–149. The two constructions of the person-based intuition are prized apart in part 2.10 below and Appendix C.

as very often *morally hazardous*. According to Variabilism, such an addition cannot, on its own, make things better but it can make them worse: it can convert an act that would be deemed permissible under an otherwise plausible permissibility theory into an act that is wrong. Variabilism reaches a different result in virtue of the fact that it considers any loss the additional person incurs to have full moral significance provided that the person does or will exist at the world at which the loss is incurred. According to Variabilism, then, any loss that the person incurs at a given world at which that person has been "added" – the loss that person incurs, not relative to the world where that person never exists at all, but rather relative to some third world at which that person enjoys a *better* existence – would count against the choice to bring that person into existence to begin with. Variabilism, accordingly, is able to avoid the objections that force us to reject both the Neutrality Intuition and the Prior Existence View. And, as we shall see, Variabilism along the way generates an interpretation of the person-based intuition that retains the full force of that intuition but also makes the intuition far more plausible than the interpretation that is suggested by either the Neutrality Intuition or the Prior Existence View.

The view that Singer calls the Prior Existence View expresses an "intuitive judgment" that Singer suggests that many people probably have – and a judgment that Totalism, the main alternative view that Singer is at that point considering, must reject. According to that judgment, the ordinary couple, contemplating whether they ought to have children, need not take into account in making their decision the "likely future pleasure of their children." Their choice whether to have children or not should instead be made on the basis of the effects that choice will have on people who do or will exist *independently* of that particular choice.[8]

Singer argues, however, that still other cases – cases in which we anticipate not that the child will be happy but rather that the child will be miserable if he or she exists at all – show that the Prior Existence View fails. Singer's argument is persuasive. We must reject the Prior Existence View. However, for purposes here, what is important is that rejecting the Prior Existence View does not mean that we are bound, after all, to say that the "likely future pleasure" of the child must be taken into account in determining whether the couple is obligated to bring that child into existence.

Variabilism makes that point. It offers an analysis that explains just why the child's prospective misery must be counted against the choice to bring the one child into existence, but also why the child's prospective happiness need not be counted in favor of the choice to bring the distinct child into existence. The upshot is that the fact that we must reject the Prior Existence View does not mean that we must reject Variabilism as well – or that we are in the end compelled to agree that the loss incurred by the merely possible person when we leave that person out of existence altogether can make the otherwise permissible act wrong.

Singer's discussion of the Prior Existence View takes him into the very deep waters of what is called the *Asymmetry*. The two claims that constitute the Asymmetry are just these. The fact that a possible person's life would be well worth living does not count in favor of bringing that person into existence and cannot, on its own, mean that it is wrong not to do that. *But* the fact that a person's life would

[8] Singer (1999), pp. 103–104.

be completely miserable – *less* than worth living – counts against bringing that person into existence – and may well mean that it would be wrong to bring that person into existence.[9] Totalism, implausibly, denies the first of these two claims. That is what makes the Prior Existence View attractive to begin with. But, at least as implausibly, the Prior Existence View denies the second of these two claims. It is for that reason that we must reject the Prior Existence View.

Having rejected the Prior Existence View, we may feel that we are forced as well to reject the Asymmetry itself. We may feel we have no choice but to consider the additional person's prospective happiness – in my terms, the *loss* the merely possible person will incur if that person is left out of existence altogether – as *morally significant* for purposes of evaluating the choice whether to bring that person into existence or not.

Variabilism, however, shows that that is not the case. According to Variabilism, what is important is *where* the losses are incurred in relation to the person who incurs them. Since, for the former claim, the loss is incurred at a world where the person who incurs that loss never exists at all, that loss has no moral significance whatsoever. So it does not count against the couple's choice to impose that loss – their choice, that is, not to have children. In contrast, for the latter claim, the loss we are worried about – the loss the person incurs having been brought into an existence less than worth having – is incurred at a world where the person does, or will, exist. So that loss, according to Variabilism, has full moral significance: it counts against the choice to bring that person into existence.

Thus, Variabilism, by taking, in a perfectly principled way, one result from Totalism and one result from the Prior Existence View, comes up with a treatment of the Asymmetry that itself seems highly plausible.

1.2.3 In Chapters 3 and 4, I explore a handful of theories that aim to throw – or, even if they do not *aim*, still, if correct, and when put together with other principles and deontic axioms we may be hesitant to reject, *succeed* in throwing – the permissibility of early abortion into serious doubt. Those theories weigh in, to one degree or another, *against* early abortion. And each does so on the basis of its own separate and, I argue, mistaken conception of just how merely possible persons matter morally.

In Chapter 3, I examine Jeff McMahan's Pareto-inspired paradox of abortion and prenatal injury (what I will call the *Abortion Paradox*).[10] McMahan argues that, for a class of scenarios that on the face of things seem morally indistinguishable, some abortions are wrong and others are perfectly permissible. What is of particular interest here is the half of the paradox that argues *against* abortion.

McMahan himself proposes to preserve both halves of the paradox. And he proposes to do so by taking note of a distinction in the *order of presentation of alternatives to the agent*.

But we may not see *why* the *order of presentation* – in contrast with, for example, the unordered contents of the relevant *feasible set* – can make such a difference to permissibility. If the logic that McMahan appeals to functions as intended in one

[9] McMahan (2009), p. 49.

[10] See McMahan (2006) and Chapter 3 below.

case, it may seem that, whatever the order of presentation, that same logic should function in the same way in the second case as well. It may seem that we end up with the result that the early abortion is itself *wrong across the board*.

I argue that a more productive approach is to challenge the Abortion Paradox itself. In particular, I question McMahan's application of what may seem to be a perfectly plausible Pareto principle. I argue that this principle – which makes an appearance not just in the half of McMahan's paradox that argues that abortion is wrong but also in the *Mere Addition Paradox* – is in fact highly questionable.

Variabilism is useful in developing this critique. Where the abortion under discussion is an early abortion, Variabilism challenges the inference from (1) the wrongness of the choice to *injure* the fetus in a case where the fetus survives and the person it develops into is then itself injured as well, to (2) the wrongness of the choice not to bring that same person into existence to begin with. The inference may seem plausible, since from the perspective of the fetus and the potential person the second loss is deeper than the first. If imposing the first is wrong, surely imposing the second is wrong as well. Appealing to Variabilism, however, we can point out that the depth of a loss is one thing, and its moral significance another. And we can argue that the loss that is incurred by the person as a result of the fetal injury has *full moral significance*; it *counts against* the choice of fetal injury. But we can also argue that the loss incurred by the person who is never brought into existence to begin with is *devoid of moral significance* – and that loss, accordingly, cannot be counted against what has been done.[11]

1.2.4 Chapter 4 examines three additional positions – from R.M. Hare, Don Marquis and Elizabeth Harman – that, whether by design or not, place the permissibility of the early abortion in doubt. Each of the three positions seems to assign at least some moral significance to the loss incurred by a merely possible person when that person is left out of existence altogether. Those positions are, in turn, at odds with the account of the moral significance of loss that Variabilism itself suggests.

Hare – like anyone who thinks we ought to *maximize wellbeing in the aggregate*; like any Totalist – can be understood to say that the merely possible matter morally in a certain robust way. Hare's view can be put in terms of *loss*. According to him, all losses incurred by all merely possible persons, including the losses they incur when we fail to bring them into existence to begin with, have full moral significance for purposes of evaluating the acts that impose those losses and their alternatives. These moral data in hand, we can then expect the otherwise plausible permissibility theory to

[11] McMahan sees the later person as *numerically identical* to the earlier fetus, the fetus that is the subject of the hypothetical fetal injury. I adopt a distinct metaphysics for purposes here. The fetus, a human organism, may well be identical to the *human organism* that develops in utero, is born and then dies. On the way I am looking at things here, however, it's not identical to the *person*, a being whose coming into existence is (I suggest) signaled by its *first thought*. This way of looking at things, I believe, lends clarity to the discussion and has an independent plausibility. But whether I am correct or not on that point, it would be a mistake to think that the *metaphysics* is ultimately what is at issue here. A different metaphysics is not going to erase the substantive moral issues that are (I will argue) raised by McMahan's Paretian analysis. At most, it would force us to *rearticulate*, not to *concede*, the underlying moral issues. See notes 24 and 143 below.

conclude in any ordinary case that early abortion – and non-conception, and conception followed by non-fruition – is wrong. That is so, even taking into account that, according to Totalism, the loss incurred by the person who never exists at all as a result of the early abortion must be balanced against the losses incurred by the pregnant woman or others if the pregnancy is permitted to proceed – even when, that is, the calculation of aggregate wellbeing is done, as it should be under Totalism, *on a net basis*.[12]

Marquis whittles down the class of those who matter morally by embracing what we can call an *existence condition*. According to Marquis, we are to worry about the loss that is imposed on the person who never exists at all *only if* the human embryo, or early human fetus, that may ultimately develop into that person *in fact exists* and has a future of a certain kind – a "future like ours" or a "future of value."[13] Thus, where Hare seems committed to the view that non-conception is just as morally problematic as early abortion, Marquis's approach targets just abortion (early and late). On the other hand, since Marquis is no utilitarian – for him, it may well be wrong to deprive the merely possible person of its future of value by way of early abortion even if that early abortion happens to maximize aggregate wellbeing – his stand against early abortion is in this way more sweeping than Hare's.

Despite their differences, however, it is clear that Hare and Marquis have in common extravagantly generous attitudes in respect of the merely possible. For both, the losses incurred by persons who, by dint of the early abortion, *will never exist at all* have moral significance for purposes of evaluating the early abortion. But Inclusion, I will argue, whether restricted in the way that Marquis suggests or not is a mistake. It is a mistake to say that the losses imposed on merely possible persons as a result of our choice not to bring them into existence to begin with have moral significance. It's a mistake to assign to the merely possible any sort of *necessary* moral status – a mistake to say that *all* of their losses have moral significance. And it's a mistake that Variabilism tries to correct, by understanding that some losses incurred by the merely possible have full moral significance and others have none at all.

Elizabeth Harman aims to reconcile two positions. The first is that early abortion is permissible even when continuing the pregnancy would impose no burden on the pregnant woman herself or anyone else. The second is that some early fetuses but not others have a certain moral status, a status that gives agents a moral reason to bring it about that those fetuses but not others develop into full-fledged persons and that makes those fetuses but not others appropriate objects of their progenitors' love. The difference, for Harman, is between the fetus that – because the agents choose against early abortion – will develop into a person and the fetus that – because the agents choose in favor of early abortion – will not. In one case, in other words, we are dealing with a person who does not yet but will exist, and in the other case we aren't.

Thus Harman, like Marquis, adopts a certain *existence condition*. But unlike Marquis's, Harman's is a forward-looking existence condition. It is a condition that

[12] See Hare (1973) and part 4.3 below.

[13] See Marquis (2008) and part 4.4 below.

is satisfied not by the mere existence of the *human embryo* or *early fetus* but rather only by the *eventual* coming into existence of the *person*. The *future person* has moral status, and hence *so does the existing fetus*. The *merely possible* person has no moral status, and hence *neither does the existing fetus*.[14]

The puzzle that is raised by Harman's approach is whether we must now say that the early abortion in fact performed is *permissible* but also that, at an alternate world where the agent chooses to continue the pregnancy, that choice becomes *obligatory*. This combination of results creates a serious conceptual difficulty. After all, if the early abortion is itself *permissible*, it seems that its alternative – continuing the pregnancy – cannot at the same time be *obligatory*.

Harman's view, by its own terms, seems to avoid any such result. On her view, the person who never exists at all in one world, where the early abortion is in fact performed, nonetheless has a certain *moral status* in an *alternate* world where the pregnancy is continued and the person eventually exists. Accordingly, agents at that alternate world have a *moral reason* to bring that person into an existence worth having. But from that fact we cannot, according to Harman, conclude that agents are there *obligated* to bring that person into existence. That is so, even in the case where agents have no moral reason not to bring that person into existence and no moral reason to do anything other than continue the pregnancy.

In short, by insisting on a certain disconnect between moral status and moral reason on the one hand, and moral obligation on the other, Harman avoids the conceptual problem.

It seems, however, that the proposed disconnect raises questions of its own. If we – unlike Harman – have a hard time prizing apart moral reason and moral obligation – if we accept that the fact that we have a moral reason to do something, and no moral reason not to do that thing, and no moral reason to do anything else, then we are obligated to do that thing – we are going to find in Harman's analysis either an inconsistency or, at best, a subtle argument *against* early abortion.

The better view, I suggest, is to begin with the view that the agent does not have a moral reason to continue the pregnancy – even at the world where that is exactly what the agent does, and where we agree that any losses incurred at that world by the person who will exist as a result of that choice will have full moral significance. Variabilism gives us a foundation for that position. According to Variabilism, the losses that have moral significance include just the losses that that person incurs at that world or any other where that person does or will exist. That means that the loss that person incurs at the world where the early abortion is in fact performed remains *devoid of moral significance*, even when the act we happen to want to evaluate – continuing the pregnancy – is performed at a world where that person does or will exist. We are thus left, according to Variabilism, with *no moral reason to explain away* – no moral reason that we must then demonstrate *not* to be connected in a natural way to our assessments of moral obligation.

The upshot is that each of these four theories, whether by the intention of their authors or otherwise, creates serious questions regarding the permissibility of the

[14] See Harman (2000) and part 4.5 below.

early abortion. Moreover, in each case, those questions arise out of the positions the particular philosopher seems to want to maintain regarding the moral significance of the merely possible person.

My own view is that we will obtain a more plausible account of early abortion when we have in hand a more plausible account of the moral significance of merely possible persons. And I think that Variabilism is just such an account. Variabilism proposes a *middle ground* between Inclusion and Exclusion. And that *middle ground on the merely possible* opens the door to a *middle ground on abortion*. Variabilism shows that we can *both* accept that the merely possible person has a moral status that is identical to our own – that is, a moral status *of sorts – and* take the position that the loss incurred by the merely possible person when we leave that person out of existence altogether has *no moral significance whatsoever* for purposes of evaluating the act that imposes that loss or any of that act's alternatives. The trick is to recognize the *variable moral status of each merely possible person – and each of the rest of us as well*.

1.2.5 The focus of Chapter 5 is abortion. I argue, with reference to Variabilism, that the choice of early abortion is ordinarily permissible and the choice of late abortion is ordinarily wrong.

How these results regarding moral permissibility are to be translated into the law – what they mean, if anything, for the law – is an important question. Both the equal protection clause of the Fourteenth Amendment of the U.S. Constitution and the due process clause are – as we shall see – critical in connection with this discussion.

Someone recently pointed out to me that what I call a *middle ground* on abortion is in fact going to offend practically everyone.[15] But I don't think it should. Those who disdain abortion at any stage of pregnancy for any reason are surely most concerned to see abortion recognized as impermissible late in pregnancy. And those who think abortion is permissible throughout the pregnancy are surely most concerned to have that fact recognized in connection with early abortion. My proposal reflects both those concerns. It asks each of the two camps to *take the ground that is most important to it – to settle for that – and concede the ground that is most important to the opposition*.

The discussion of early abortion may well remain fraught. But that may be less moral law and more a matter of evolution: a principle of genetic survival that leads many men – and no doubt many *grandmothers* as well, and let's not even talk about *grandfathers* – to find in moral law edicts commanding women to have babies. The discussion of late abortion may remain fraught as well. But there, too, we can identify a source of angst beyond moral law. We can recognize that there is something very wrong with a community that requires women to do so much for their soon-to-be-born offspring and allows everyone else to do so little for anyone at all. And we may well wonder whether one person's choice can be deemed wrong when everyone else is doing things that are at least as bad and

[15]I owe this point to Mark K. Greene.

getting away with them. In the end, however, it seems that these ancillary facts are not going to change the conclusion that we ultimately reach here – that the choice of late abortion is ordinarily wrong. That there's something wrong with the *community* does not imply anything at all about what is permissible for the pregnant women or any other *individual* agent to do.[16]

1.3 Inclusion, Exclusion and a Dilemma

1.3.1 Does the loss a merely possible person incurs when we leave that person out of existence altogether *count against* the choice that imposes that loss? Can that loss make the otherwise permissible choice that imposes that loss wrong?

We can clarify this question. For any particular way the future might unfold – that is, for any possible future or *world* – we can define a class of *merely possible persons*. Merely possible persons are persons agents *could* but in fact *don't* bring into existence at that world. The choice not to conceive a child, or to conceive a child in vitro but not arrange for the newly created embryo to develop *in utero* into a thinking thing, a *person*, and (by assumption) the choice to have an *early abortion* are all choices that relegate, relative to the world of their performance, *possible* persons to the class of the *merely* possible.

Suppose that agents *do not make* but *could have made* the future unfold for a particular person in such a way that that person comes into existence. And suppose that those same agents *could have made* the future unfold for that person in such a way that that existence itself is *worth having*. Suppose, in other words, agents had the alternative of bringing a particular person into existence *without* dooming that person to either a *wrongful life* – a life *less* than worth living – or a *perfectly neutral* life – a life whose burdens *perfectly counterbalance* whatever it is that is precious in life to the one who lives. And suppose, finally, that we are not in a case where the choice *not* to bring a person into an existence worth having makes things better for anyone at all. This not a case where agents, for example, *must* leave one person

[16] The argument against wrongdoing that is proposed here may sound vaguely like the problem of *collective harm*. The problem of collective harm is urgent. Where, *whatever* I as an agent do, others will suffer by virtue of what still other agents will do, it may seem unclear whether what *I* do can harm those others. And it, accordingly, may seem unclear what my obligations in that context really are. I argue elsewhere that the problem of collective harm does not absolve us from our *participation* in the group act that leads to that harm. I can, as a *participant in a group*, impose harm even if I have not, as an *individual*, imposed harm. The problem of collective harm arises in virtue of the fact that we have overlapping agents: groups of individual agents, overlapping with individual agents. The *group* may have a way to create more wellbeing for a person rather than less when the *individual* does not. Recognizing that fact can help us, I argue elsewhere, to address the problem of collective harm. See Roberts (2006). But the problem presented by late abortion is quite different. There, the question is whether, if everyone else is blithely doing wrong things, why is it wrong for me to do wrong things as well. And the answer is: it just is; it is wrong to do wrong things.

out of existence altogether in order to avoid making things much worse for many others.

Must we then say that the choice not to bring a person into an existence worth having is *wrong*? Must we at least say that agents have a *moral reason* to choose otherwise – or that, *other things being equal*, or *ceteris paribus*, that choice is wrong, or that it is *prima facie* wrong?

After all, there is at least a case to be made – and I think an exactly correct case – that the choice not to bring a new person into an existence worth having is a way of creating less wellbeing for that person when agents could have created more – that that choice makes things worse for a person when agents could have made them better. There is a case to be made, in other words, that the choice not to bring a person into an existence worth having causes that person to incur – though obviously not to *suffer* – a *harm*, or *loss*.

Given, then, that the loss is *incurred*, does it *count*? Does that loss *bear on* the permissibility of the choice that imposes that loss or the alternatives that avoid it? Is that loss *morally significant* for purposes of evaluating that choice? Does it *count against* that choice? Does that loss, in a roundabout way, count *in favor of* the choices that avoid it? If the choice itself is otherwise permissible, *can that loss make that choice wrong*?

1.3.2 Let's first consider the claim that, yes, the loss under scrutiny can make the otherwise permissible choice that imposes that loss wrong. This claim may seem plausible on its face. Some theorists might even consider it obvious that, when we don't have a good reason *not* to bring a person into an existence worth having, we are *obligated* to bring that person into that existence.

Theorists who accept this claim seem to be adopting a radically *inclusive* position in respect of the merely possible. They may be thinking that the *merely possible matter morally*. They may be thinking that we have *obligations in respect of the merely possible*. And they may be concluding that all of the losses incurred by the merely possible, *including* the losses they incur when we choose not to bring them into existence to begin with, are *morally significant* for purposes of evaluating the acts that impose those losses and the alternatives to those acts that avoid those losses.

Now, there seem to be two basic ways of grounding this position. The first rationale asserts that we have obligations in respect of merely possible persons as individuals, one-by-one *for their own sake*. Such obligations can be described as *person-affecting*, or *person-based*, in nature: they are obligations to *affect* (under certain conditions) *particular persons for the better and not for the worse*. Since one way of affecting persons for the better seems to be to bring them into existences worth having in place of none at all, the argument can be made that this is something we are morally required to do – at least when we can do it without making things too much worse for too many others.

To say that we have obligations *in respect of* the merely possible is not necessarily to say that those obligations are *owed to* anyone at all. Still, talk about our having obligations in respect of *merely possible* persons as individuals, one-by-one for their own sake, may seem peculiar. And certainly the *specific* implication that leaving a person out of an existence worth having could ever on its own ground a finding

of wrongdoing is ruled out by the narrow person-based principles I have elsewhere proposed.[17]

But even if talk about our having obligations in respect of *merely possible* persons as individuals, one-by-one for their own sake, seems peculiar, the idea that we have such obligations, and that we are accordingly obligated, in many cases, to

[17] The person-based principles I have elsewhere described include the *person-affecting*, or *person-based*, intuition. That core intuition provides just a *necessary*, not a *sufficient*, condition on wrongdoing. The particular version of the person-based intuition I elsewhere adopt – and that is part of Variabilism in combination with an otherwise plausible permissibility theory and specifically OPPP2; see part 2.2.4 below – is the traditional narrow version. According to that principle, an act performed at a world is wrong *only if* that act makes things worse for at least one person *who does or will exist* at that world than things are for that same person at some alternate world accessible to the agents. The act that, performed at a given world, is *bad* must be *bad for* – must, that is, *harm*, or impose a *loss* on – some person who does or will exist at that world. Put in the terms of Variabilism: that act must impose a *morally significant loss* on a person.

Under this narrow version of the person-based intuition, an act that leaves a person out of an existence worth having cannot be declared wrong just on the ground that it imposes a loss on that *merely possible person*. Again, in terms of Variabilism, the loss that act imposes is not itself *morally significant*. If such an act is wrong at all, it is wrong on other grounds.

A "wide" version of the person-based intuition, in contrast, might assert that the loss incurred when a merely possible person is left out of an existence worth having can on its own ground a finding of wrongdoing. This version is represented by Inclusion in combination with OPPP2; see part 2.2.4 below. The wide version of the intuition is not very interesting. (We may even think it misses *the intuition* altogether.) In any case, OPPP2 itself is consistent with both the wide and the narrow version of the person-based intuition.

At the same time, the version of the person-based intuition I describe elsewhere and adopt here (as an implication of Variabilism in combination with OPPP2) is broader in a certain critical respect than what John Broome calls the *Neutrality Intuition*. Broome (2004), p. 143; see also part 2.10 below and Appendix C. On the version I adopt, the necessary condition on wrongdoing is satisfied whenever agents create less wellbeing for an existing or future person when they *could have created more* for that same person. According to this view, when we compare two acts in order to determine their respective *moral betterness*, and ultimately the wrongness of one of those two acts, the loss that counts toward moral *worseness*, and ultimately wrongness – the loss that counts as *morally significant* – can be established, not only by our noting that the one act is worse for an existing or future person than the second act is, but also by our noting that the one act is worse for that person than *some third act*. The Neutrality Intuition operates quite differently, and Broome – rightly, I think – rejects it. The Neutrality Intuition itself, along with the related Prior Existence View that Singer earlier articulated and rejected (Singer 1999, pp. 102ff.), is discussed in part 2.10 below. Broome's more sweeping objection against the person-based approach generally is discussed in Appendix C below.

Another clarification: the person-based approach I have elsewhere described is not equivalent to Exclusion (or Moral Actualism); my principles (barring an occasional lapse; see Roberts (2002), p. 328, where it should have been made explicit that each person who "exists in Y also exists in X") do not leave the merely possible and all their losses out of the picture altogether. They do not, that is, imply that merely possible persons *cannot* be wronged – or, more specifically, that the losses they incur *cannot* make the acts that impose those losses wrong and *cannot* make the acts that avoid those losses right. My principles by their terms apply to all acts performed at all worlds, actual or not. And they are, at the same time, restricted to reflect the point that an act that makes things worse for, for example, *actual* persons is not necessarily wrong in a case where the only way of making things better for those persons would have been to bring still other persons into existences that are less than they might have been. See Roberts (1998), (2002), (2003a), (2003b), (2004), (2006), (2007), (2009).

create more wellbeing for them rather than less, including by way of bringing them into existences worth having, is something we seem able to make sense of.

A second rationale may sound less peculiar but just as effectively ground the idea that we have obligations in respect of the merely possible and that all of the losses they incur have moral significance for purposes of deciding what we ought to do. This second rationale is *impersonal* in nature. It states that a very important moral obligation, and perhaps our only moral obligation, is to *make things better*, not for any person, existing, future or merely possible, but rather *better simpliciter*. We are obligated, in Sidgwick's phrase, to make things better "from the point of view of the universe."[18]

This second rationale, like the first, asserts that we have obligations in respect of the merely possible. But, more clearly than the first, it does not require that we say that those obligations are *owed to* anyone at all. We can be obliged to make certain choices – choices that create additional wellbeing on an aggregate basis, or choices that create additional quantities of the *overall good*, which might itself be an amalgam that includes, in addition to aggregate wellbeing, goods of other sorts as well (such as equality). On this second rationale, we have such obligations not because those choices make things better *for any person*, existing, future or merely possible, but rather because those choices *make things better simpliciter*. We can then say that bringing people into existences worth having is often a way of doing just that.

The claim, then, that the loss incurred by a person we never bring into existence to begin with can make the otherwise permissible choice that imposes that loss wrong is one we can explain. We can say that the merely possible *matter morally*, and that *we have obligations in respect of them*, either one-by-one for their own sake, or on impersonal grounds. And we can say that, accordingly, the losses they incur – *including* the losses they incur when we leave them out of existence altogether – *count against* the otherwise permissible acts that impose those losses. Their losses can make those acts *wrong*.

1.3.3 The other of the two possible answers to our initial question is just that, no, the loss incurred by a merely possible person when we leave that person out of existence altogether *cannot* make the otherwise permissible choice that imposes that loss wrong.

Variabilism enables me to make this point more clearly. Thus: a person who counts as merely possible relative to a world w but who does or will exist relative to a world w' can incur a morally significant loss at w' – a loss, that is, that *bears on* the moral status, not just of acts that *impose that loss at* w', but also of acts that *avoid that loss at* w.

The critical point is that Variabilism, in combination with the otherwise plausible permissibility theory, clearly establishes that a person-based approach is *not* equivalent to Moral Actualism. A person-based approach *need not*, even if some do, put *actual* persons – or even persons who *would be actual* were the particular act under scrutiny itself performed – on a moral pedestal. Among other things, this means that, even though Caspar Hare's arguments against Moral Actualism were not mistaken, his conflation of Moral Actualism with the person-based approach was. Hare (2007), n. 5.

[18] Sidgwick (2006/1884), Preface.

Theorists who make this claim seem to be adopting a radically *exclusive* position in respect of the merely possible. They may be thinking that we have no moral obligations in respect of merely possible persons at all – or, at least, that the merely possible do not matter morally in the same way that you and I do. How could they? They don't, and won't, exist. It might accordingly seem that none of their losses, including the losses they incur when we fail to bring them into existence to begin with, have any moral significance whatsoever.

Suppose that a young woman has chosen non-conception, or conception followed by non-fruition, or early abortion. We agree, of course, that the young woman matters morally. So does her partner and her already-existing and even her future children, and no doubt her grandmother and many others as well. We think their losses are morally significant and must be taken into account in determining whether what the young woman has done is permissible. We think their losses, if any, bear on the permissibility of her choice of early abortion. But we might deny that the losses incurred by merely possible persons – including the loss incurred by the child the woman never brings into existence to begin with – have any moral significance at all.

This is not to deny that there's a *loss*. It's just to say that that loss cannot make the otherwise permissible act that imposes that loss *wrong*.

1.3.4 Our initial question was this: does the loss a merely possible person incurs when we leave that person out of existence altogether count against the choice that imposes that loss? Does such a loss make the otherwise permissible choice that imposes it *wrong*?

We now have two answers to that question – one that claims that there *is* something wrong with the choice not to bring the merely possible person into an existence worth having in place of none at all, and one that claims that there *isn't*. And we can provide rationales for each of those positions. We can adopt Inclusion and say that we *do* have obligations *in respect of* (though not necessarily *owed to*) the merely possible, whether, one-by-one for their own sake or on impersonal grounds, for the sake of the universe, and that their losses accordingly can make the otherwise permissible acts that impose those losses wrong. Or we can adopt Exclusion, and say that we *don't* have obligations in respect of the merely possible and that the losses incurred by the merely possible accordingly *cannot* make the otherwise permissible acts that impose those losses wrong. Under Inclusion, the merely possible *matter morally*, and *so do their losses*, including the loss of never coming into existence at all. And under Exclusion they *don't* matter morally and *neither do their losses* – especially not the loss of never coming into existence at all.

1.3.5 The position that *includes* the merely possible among those who matter morally and deems all their losses to have full moral significance – *Inclusion* – seems at odds with even the most basic conception of procreative privacy. For Inclusion appears to ground a powerful argument against non-conception, against conception followed by non-fruition *and* against early abortion. At least, Inclusion vastly complicates what we think should be a summary moral analysis.

Suppose that a young woman chooses to have an early abortion when she could have instead continued her pregnancy and brought a new person – the person who

is relegated by her choice to the class of the merely possible – into an existence worth having. And suppose that Inclusion is true – that merely possible persons matter morally and their losses count. In evaluating her choice, we are then compelled to take into account *not just* the losses incurred by or avoided on behalf of the woman herself and each other person we all agree matters morally – her partner, her already-existing children, her future children, her grandmother, etc. – but also the loss incurred by the merely possible child. And we must recognize that the loss incurred by that child is very great and cuts very deep: the woman's choice of early abortion leaves her possible child *with nothing at all*. In contrast, the woman herself and each other person we agree matters morally may well have already accrued a good deal of wellbeing over their lifetimes and may well reasonably expect to accrue a good deal more whether the early abortion is performed or not. Let's suppose that this is in fact the case. Trading off, then, the various losses against each other – the pregnant woman's, for example, if she proceeds with the pregnancy, against her possible child's, if she does not – we seem compelled to say that the woman's choice of early abortion – or non-conception or conception followed by non-fruition – is *wrong*.

1.3.6 But this result seems completely untenable. We seem here to have generated an *implausibly stringent procreation obligation* – and not just on the part of young women. Surely we must reverse course and say that merely possible persons, after all, *do not matter morally* – that *existing* persons matter morally, and *future* persons matter morally, but the merely possible matter morally *not at all*.

The position that the merely possible are to be *excluded* from those who matter morally and that their losses are, accordingly, devoid of moral significance – *Exclusion* – seems, at first glance, the far more commonsensical position. Isn't it just obvious that what is important is for us to "make people happy" – that "making happy people" cannot really be on our moral to-do list[19]? Surely it is enough that moral law demands that we see to the needs and interests of all those persons who do or ever will exist. But it seems too much – a perversion rather than an expression of moral law – to think that we must also see to the needs and interests of all those many, *many* persons who will never exist at all.

1.3.7 What makes things interesting, however, is that we *can't* reverse course completely without finding ourselves in, if anything, a still *more* untenable position. For, if nothing else, we should be able to agree that distinct persons exist in distinct worlds. *We* need *not* have existed; persons other than ourselves *might* have existed. This means – as we shall see – that it is not going to work to say existing and future persons matter morally but the merely possible matter do not. It is not going to work to say that *we*, at *this* world, matter morally and *they*, at *that other* world, do not.

Now – as we shall also see – Exclusion can itself be developed in either of two ways, *Alpha* and *Beta*. But Exclusion in both its forms collides with the fact that our assessments of what goes on at the distinct worlds and within their distinct

[19] I am paraphrasing Narveson (1976), p. 73.

populations are closely linked to each other. The upshot is that the very same act, under Exclusion Alpha, will be deemed obligatory if unperformed and wrong if performed, a result that unfathomably assigns a moral significance to the fact of performance, and, under Exclusion Beta, will be both permissible and wrong in the very same case – a result that seems clearly inconsistent. To put the point – for now – roughly: the idea that existing and future persons matter morally and the merely possible do not is going to come at the cost of any stable sense of just what we are trying to say when we say that one act is permissible and another is not.

1.3.8 So we have a dilemma: a completely untenable position if we include the merely possible among those persons who matter morally and say that all their losses have moral significance for purposes of determining what we ought to do, and a still more untenable position if we exclude the merely possible, and say instead that none of their losses have moral significance for purposes of determining what we ought to do. Inclusion is untenable, and Exclusion, in both its forms, is, if anything, still more untenable.

1.4 Variabilism as Middle Ground

The idea that Exclusion and Inclusion are our only options – that, if we can't be Exclusionists, we must be Inclusionists – sometimes seems to drive the argument in the minds of some theorists. But it is a false dilemma. And that that's so can be seen when we take care to focus, not on who matters morally and who does not, but rather on what losses matter morally and what losses do not. We can then easily reject Exclusion in both its forms without accepting Inclusion. Instead of saying that *no* losses incurred by any merely possible person have any moral significance at all, or saying that they *all* do, we can say that *some* losses incurred by the merely possible have moral significance and *some* do not. *Exclusion* is a mistake, but so is *Inclusion*.

And we can exactly the same thing about the losses incurred by all of the rest of us as well. We can say that *we all* – existing and future persons, *and* the merely possible – matter morally *and* in exactly the same ways. We all matter morally, but we all matter variably. Now, this might come as a surprise. We might have thought that, having made it into existence, we ourselves are to be accorded some special moral status. But that's not so, according to Variabilism. Some of our losses have moral significance, and some of our losses do not.

Moreover, the variability we see in moral significance is not at all random. There is a distinct pattern. Thus:

Variabilism:

> Any loss incurred in any world by any person has moral significance for purposes of evaluating an act that imposes that loss and any of that act's alternatives *if and only if* the person who incurs that loss does or will exist at the world at which that loss is incurred.

On this view, what imbues some losses but not others with moral significance is just *where* that loss is incurred in relation to the person who incurs it. Incurred at world where the person does or will exist, the loss has full moral significance. Incurred anywhere else, the loss has no moral significance at all.

1.5 Variabilism and Abortion

1.5.1 Exclusion and Inclusion are each very extreme positions. As noted above, they accordingly leave a good deal of middle ground between them. They leave room for *Variabilism*. And that room, in turn, opens the door to a certain middle ground on abortion.

The implication from Variabilism that I will argue is important here for dealing with early abortion is just that the loss incurred by a person at a world where that person is never brought into existence to begin with has no moral significance whatsoever. And the implication that is important for dealing with late abortion is – I will argue – just that the loss incurred by a person at a world where that person does or will exist has full moral significance.

1.5.2 Specific permissibility results will need to wait on the otherwise plausible permissibility theory. That theory is spelled out in Chapter 2 below. And, as noted earlier, for purposes here, certain critical assumptions relating to abortion will be adopted. One is that the early abortion is the choice not to bring a person into existence to begin with. Another is that late abortion is the choice to end the life of an already-existing person. We will say more about these assumptions, and others, briefly in part 1.6 below and in more detail in Chapter 5 below.

Still, looking ahead, we can outline the position that Variabilism in combination with the otherwise plausible permissibility theory will take regarding early abortion. *Early abortion* is ordinarily perfectly morally permissible so long as it is what the pregnant woman wants. That is so, since the only person who will typically incur any loss – or at least any loss that is on par with the loss the woman may incur if she continues the pregnancy – as a result of the early abortion is the person the embryo or early fetus will develop into if the pregnancy proceeds and the person who never comes into existence at all (never having had that first thought) if the early abortion is performed. According to Variabilism, that loss, however great or deep, is devoid of any moral significance whatsoever. It thus cannot make the otherwise permissible choice to have the early abortion wrong.

1.5.3 Again looking ahead, we can say that the position that Variabilism in combination with the otherwise plausible permissibility theory will suggest regarding late abortion is this. *Late abortion* is ordinarily wrong. Late abortion ends the life of an already-existing person, a being that has had that first thought and acquired the capacity to survive. It thus causes that person to incur a loss that, according to Variabilism, has *full moral significance*. Moreover, that loss is typically very great and very deep, at least when compared against other losses that might be incurred by other persons if the late abortion is not performed. It, accordingly, cannot be justified under an

otherwise plausible permissibility theory. That is so, even in cases in which the woman's own health or even life is at risk or the child faces an existence that is worth having but unavoidably flawed. The exception is the unusual case where the child's life will inevitably be either wrongful or perfectly neutral. In that case, the late abortion imposes *no loss at all* on the child and hence no morally significant loss.

1.5.4 Variabilism thus suggests – for any typical case – a hard line on late abortion and almost no line at all on the early abortion. Early abortion is typically *perfectly* permissible – and a choice we should find permissible on the basis of a *very* summary moral analysis if it is what the pregnant woman herself wants. There is no point in any nuanced, probing, multifarious, dense analysis of just exactly how (A) the loss – if any – the woman may incur if the pregnancy is allowed to continue is to be traded off against (B) the very great and very deep loss her merely possible child will (we agree) incur if the early abortion is performed. In fact it is a symptom of moral confusion, not moral sensitivity, to think that determining the permissibility of the early abortion requires us to carefully balance (A) *any loss at all* (morally significant or not) against (B) a loss that is *completely devoid of moral significance*. Such balancing acts can be sensible *only if* the losses that are to be balanced against each other *both have moral significance*.

The choice of early abortion is thus not ordinarily a particularly hard case. But nor is the choice of late abortion. For there we are dealing with a loss that has *full moral significance*, according to Variabilism, in virtue of the fact that the new person who will incur that loss if the late abortion is performed already has come into existence and acquired the capacity to survive. That is so, by definition of *late abortion*. That loss is, moreover, very great and very deep – at least, compared to the losses others can expect to incur if the late abortion is not performed. Thus, even when the pregnant woman herself has much to lose if she continues the pregnancy, the late abortion typically will be wrong.

Cases involving wrongful life – or even the perfectly neutral life – are the exception.[20] In those cases, the late abortion imposes *no loss at all* on the already-existing person. If the case involves the genuinely wrongful life, the life *less* than worth living, the late abortion *avoids* a loss – and a morally significant loss, at that – on behalf of that person. Such cases may be rare. But they still need to be taken into account.

1.5.5 The construction of a middle ground is sometimes a matter of trying to make things fit together that don't. It's a way of pulling ideas together from both sides of a debate in a way that really just makes things conceptually worse – harder to understand, test and apply, and harder to situate within a broader moral theory.

[20] We all are vulnerable to having our lives made less than worth living, however well-off we might now be or might have been. *Wrongful life* is a more limited concept. The life is *wrongful* when the existence is less than worth having *and* where an existence worth having, or at least a perfectly neutral existence, is not something agents *could* have brought about. In the typical wrongful life case, the possible world where the person born with the serious genetic or chromosomal disorder – so serious that the life is made less than worth living – is given a life that is worth living or at least neutral is *inaccessible* to agents. We do not, that is, yet have the medical technology required to repair or ameliorate the underlying disorder.

But I believe that the two middle grounds I propose here, one on the moral significance of merely possible people and the other on abortion, make things better in those respects rather than worse. We are left with two positions that are both plausible in their own right and fit nicely together. They together enable us to maintain a highly intuitive position regarding abortion without having that position kicked out from under us by an idea about the merely possible that we mistakenly think we are compelled to accept. We can accept the permissibility of early abortion free and clear, in other words, once we reject the idea that the loss incurred by a person when we leave that person out of existence altogether has any moral significance at all.

1.6 Thinking Things, Persons and Abortion

1.6.1 The idea – and an assumption here – that *persons* are the kinds of thing we generally have obligations in respect of and the kinds of thing we often must make things better for *at least* if they exist is not especially contentious. Many of us think that there is at least a sense of the term *person* that implies exactly that moral weight.

Accordingly, other than perhaps Variabilism itself, the most contentious point made so far may just be that a pregnancy does not involve a person until the point in the pregnancy at which *thinking* itself has emerged. The most contentious point is that *persons* are the kinds of things whose coming into existence is *signaled by their thinking*. On this way of looking at things, the embryo and the early fetus are live human organisms. But, until thinking has emerged, there is no person there.

By implication, then, the abortion that destroys a human embryo or early fetus that has not yet had that first thought does not impose a loss on an already-existing person but rather thwarts the coming into existence of a new person.

This account of when something is a person – when a thing is the kind of thing that matters morally – could be called the *Thinking Thing Account*. While I believe that account is plausible and will say some things in its favor in what follows, for purposes here it functions as an *assumption*.

Of course, there are competing views. What we can call the *Human Organism Account*, for example, asserts that human embryos and early human fetuses all count as *persons*. For purposes here, however, that account is, as a matter of assumption, false.

1.6.2 A question, however, immediately arises. Whatever our views about the obligations we have in respect of *merely possible* persons, we can surely agree that ending an *existing* person's life is morally problematic and often wrong. By making the assumption that the human embryo and early fetus are *not* persons – by accepting the Thinking Thing Account as an assumption – I escape the immediate implication from the Human Organism Account that the early abortion ends the life of an existing person and is, accordingly, often wrong. So haven't I begged an important question regarding the moral permissibility of the early abortion?

No. A question for this book is the moral significance of merely possible persons. A second, closely related, question is whether the early abortion is wrong *in virtue of what it does to the merely possible person*. But what is *not* a question for this book is whether the early abortion is wrong *in virtue of what it does to the*

already-existing human embryo or early human fetus. I am, accordingly, begging no question of interest here.

For purposes here, then, *if* the early abortion is wrong at all, it is wrong *not* in virtue of what it does to the already-existing human embryo or early human fetus but *rather* in virtue of what it does to the *merely possible person* – the person whose coming into and continuing in existence is prevented by the early abortion.

1.6.3 It's true that the abortion theories I examine in what follows, when put together with certain deontic axioms and other principles we may find hard to reject, build cases against the permissibility of early abortion in many cases in which I will argue in favor of the permissibility of early abortion. At the same time, none of the theories I examine – the theories, that is, from McMahan, Hare, Marquis and Harman – are rooted in the idea that the human embryo or early human fetus – that is, the live human organism that has not yet has its first thought – is a person. Rather, to the extent that those theories put pressure on us to conclude that the early abortion is often wrong or at least morally suspect, they do so on the ground that the human embryo or early fetus can be correlated with a certain *possible person* – a certain person who shall remain *merely possible* if the early abortion is performed.

In other words: according to the theories I will be examining, it is what we do to the *merely possible person*, not what we do to the *human embryo* or *early fetus*, that makes the early abortion wrong in any case in which it *is* wrong.[21] In adopting the Thinking Thing Account, I am accordingly begging no question that is at issue in this book.

1.6.4 While the Thinking Thing Account is an assumption for purposes here, it is worth noting that it is a highly plausible assumption. We need only look at the issues in first-person terms. *I* matter morally. And *I'm* a thinking thing; at least, I *intermittently*[22] think. But I don't think *I* will continue to exist after I have whatever last thought it is I will ever have. It is perfectly possible that *I* will go out of existence even though (however unfortunate for others) my body lives on for awhile, its vital signs strong.

But we should say about the beginning of life what we say about the end: just as *I* may cease to exist even though my body survives, so was there a time when *I* had not yet come into existence even though my body had. The same is true of *you*. We, unlike our bodies, are the kinds of things whose coming into existence and going out of existence are signaled by our *thinking*. And we, moreover, are surely among the kinds of things that *matter morally*, if variably.

---

[21] Thus Earl Conee notes that the following argument can seem sensible to any theorist who thinks that an act's permissibility is determined, at least in part, by the consequences of that act: "Assuming that an early fetus is not a person, consequentialist considerations still argue against the moral permissibility of some early abortions. If an early fetus would grow to become someone who would lead a sufficiently valuable existence, then on objective consequentialist grounds it would be seriously wrong to kill the fetus." See Conee (1999), p. 629.

[22] On the view I describe, once in existence, a person can continue to exist even if that person does not continuously think. *Intermittent* thinking alone is enough to keep one in existence. The person doesn't cease to exist, then, until the conclusion of that last thought – even if their bodies, sustained by natural or artificial means, manage to cling to life far beyond that point. In short, our first thought signals our coming into existence at a world and our last thought signals our going out of existence at that world, but, between these two points, we *exist* whether or not we *continuously* think.

1.6.5 One other point in favor of the Thinking Thing Account can be noted. What it really says is that *thinking* and *mattering morally* – that is, being a *person* – come together. It says that thinking things and thinking things alone are the kinds of things in respect of which we can have moral obligations, whether one-by-one for their own sake, or on impersonal grounds, as a matter of making things better *simpliciter*.

Perhaps what makes a connection between thinking and mattering morally most plausible is just that thinking things are such that it matters *to them* (whether they think so or not) how we treat them. It matters *to them*, that is, whether we create more wellbeing for them rather than less.

Of course, the concept of *thinking* can be naturally analyzed in many different ways. If it is correct that the reason that some individuals matter morally is that it matters *to them* whether we create more wellbeing for them rather than less, then this fact can serve as a guide to just how narrow or broad the concept of thinking we will want to adopt for purposes here needs to be.

Thus, the pain I feel when vaccinated is enough to show that I am a thinking thing for purposes here – and that I say "Not this *again*" enough to show I have continued as the *same* thinking thing from one moment to the next. But individuals do not need a life plan and or any clear concept of themselves as continuing beings to have these kinds of experiences.[23] It can *matter to them* what we do to them, just as it *matters to me* what is done to me. Moreover, we can take this approach without limiting the concept of wellbeing to simple pleasure. Thus, it can matter to the 10-year-old child whether his parents are saving for college or not, and the parents' saving for college may well create more wellbeing for that child rather than less even if the child himself is completely oblivious to the whole thing. We do not need ever to know that wellbeing is being created for us in order for wellbeing to be created for us.

1.6.6 The Thinking Thing Account is, then, though just an assumption, plausible. It is plausible that, in the case of persons, thinking and coming into existence come together. Accordingly, the early abortion – the abortion that takes place during the period the live human organism, whether embryo or early fetus or even scattered object consisting of sperm and egg, has never had that *first thought* that would signal the coming into existence of a new person – is not a choice to *end the life* of an already-existing and otherwise continuing person. Rather, it's the choice *not to bring that person into existence to begin with* – or, if not that, then the choice to end the life of an existing person whose going out of existence was naturally imminent *in any event*.

1.6.7 I argued above that the Thinking Thing Account begs none of the questions I mean to address here. What would beg the question would be to insist that persons

---

[23] But we need also to make sure that the concept of thinking we make use of here – that we want to connect with being the kind of thing that matters morally – is not *too* broad. See note 3 above and parts 5.2 and 5.3 below. In my household, there has been a good bit of discussion regarding whether insects *think* in the sense that they feel pain. My son thinks they do: how otherwise could they manage not to get themselves killed the moment they are out on their own? I asked my spouse: do you think insects feel pain? His response was: what do you mean by *pain*? I explained: pain in the sense of what I feel when I say "*Ouch, that hurts*!" Do they say: "*Ouch, that hurts*!"? He replied: yes, I think that's *exactly what happens*. They say "*Ouch, that hurts*!"

who remain *merely possible* at a given world (e.g. the *actual* world), though the *kinds* of things that matter morally, in fact *do not matter morally*. It would beg the question to insist at this early stage that how we treat the *merely possible* – whether, that is, we create less wellbeing for them rather than more – has *no moral significance* for purposes of determining whether the choice of early abortion is permissible. And in fact that is nothing we insist on now.

Moreover, it is nothing we will ever insist on at all. Part 1.7 below makes that point. Variabilism thus in fact rejects the idea that the merely possible do not matter morally – that they have *no* moral status whatsoever. Variabilism instead asserts that merely possible persons do matter morally, but that they matter *variably*, with some of their losses being assigned full moral significance and others assigned no moral significance at all.

1.6.8 One other assumption that I make here bears comment. It is a metaphysical assumption, and it asserts that the person, whose coming into existence is signaled by its first thought, and the early human fetus, whose coming into existence is signaled by such things as cell differentiation and functioning biological systems, are numerically distinct. I think that this is a useful and intuitive way of looking at things. It is clarifying to be able to recognize frankly in connection with the end of life that there are some events my body undeniably can live through but *I* can't. And – on the independent assumption that *mattering morally*, that is, being a *person*, and *thinking* come together – it clarifies things, to recognize that the developing biological organism in the form of the early human fetus comes into existence before the person does.

But in this case the assumption we are dealing with is not critical to the debate. Even if the truth is that the early fetus is numerically identical to the later person – even if what I am calling a *person* is really just the *person-phase* of a continuing biological organism – that fact will not change the substance of the moral problems that we are trying to address here or the results we shall reach regarding early and late abortion. Those problems – the problems of the moral significance of merely possible persons and of abortion – would need to be *rearticulated*, but they won't go away. Nor would any such rearticulation decide for us just what the correct resolution of those problems really is.[24]

[24] Thus, if an alternate metaphysics turns out to be correct – if, in particular, it turns out that we are each non-thinking embryos and early fetuses for awhile, and then, when thinking emerges, we become persons – then the debate here will need to be rearticulated. Instead of talking about when a person comes into existence, we will talk in terms of when the live organism transitions from its *non-person-phase* to its *person-phase*. And we will need to revise the definition of *early abortion*. It will become, not the choice not to bring a person into existence, but rather the choice to prevent the live organism from transitioning to its "person-phase." And the issue of moral significance of the merely possible will become, not whether we have obligations in respect of merely possible *persons*, but rather whether we have obligations in respect of individuals who have not yet transitioned and never will transition into their respective *person-phases*. But a different metaphysics is not going to make the substantive moral debate go away. See note 143 below. For discussion of the "irrelevance of metaphysics to the moral issue," see Conee (1999), p. 620 and, more generally, pp. 619–646.

1.7 The New Abortion Debate

1.7.1 The claim that thinking signals the coming into existence of a *person* – that until there is thinking there is no person – is, for my purposes, just an assumption. That same claim is the focus of an older, but important, debate on abortion. In that older debate, we would be talking about whether the human embryo or early human fetus is itself a *person* – that is, the kind of thing that matters morally *at least* if it exists; the kind of thing in respect of which we have moral obligations; the kind of thing we normally ought not kill. And we might then extend the debate in a certain way, following Judith Thomson, and ask whether the fact that the human embryo or early fetus is a person necessarily means that abortion is wrong.[25]

The new debate on abortion takes a different tack. It has no need to consider Thomson's claim that the fact that the early embryo or fetus is a person does not necessarily mean that early abortion is wrong. For the new debate can *assume* that the human embryo and early human fetus are not persons – and can accept that the early abortion, if wrong at all, 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 person that embryo or fetus might eventually develop into. The new debate thus focuses on how we should think about the moral significance of the *merely possible person* – the person who will never exist at all in the case where the agent chooses early abortion (or non-conception, or conception followed by non-fruition) over conception and fruition.

1.7.2 It is, then, the new debate and not the old debate that I want to take part in here. The following argument reflects that new debate. In fact, in one form or another it dominates that debate. Even theorists who do not want categorically to endorse the conclusion of this argument have felt forced to qualify their own contrary conclusions in ways that leave us wondering whether they really think real-world early abortions can ever be plainly and unconditionally justified at all.

Thus, we might think that:

1. Merely possible persons *matter morally*, in just the way that you and I matter morally.

And:

2. Early abortion often imposes a very great and very deep loss on a merely possible person when the agent could have avoided imposing that loss without imposing too great or too deep a loss on others.

And from (1) and (2) we might conclude that:

3. Early abortion is *often wrong*.

Premise (1) asserts that we have moral obligations in respect of merely possible persons, whether for their own sake or on impersonal grounds, just as we do in respect of each other. The fact that we create less wellbeing for the merely possible

[25] Thomson (1971), pp. 47–66.

when we could have created more without imposing too great or too deep a loss on others, just as when agents create less wellbeing for *us* when they could have created more without imposing too great or too deep a loss on others, is, according to (1), the sort of thing that can make the otherwise permissible act wrong.

Premise (2) asserts that early abortion often imposes a very great and very deep loss on a merely possible person when the agent could have avoided that loss on behalf of that person *without* imposing too great or too deep a loss on anyone else, including the pregnant woman. The difference is that the merely possible person accrues no wellbeing at all – ends up, that is, with the zero wellbeing level implied by never having existed at all – if the early abortion is performed, while the pregnant woman typically has accrued a good deal more wellbeing than that during her lifetime and can expect to accrue a good deal more whether she has the early abortion or not.

Now, premises (1) and (2) both require further discussion, and we will come back to them shortly. But we should begin by taking a quick look at the conclusion – that is, (3). Just how plausible is (3)? Consider the case where the early abortion – the choice *not* to bring a new person into existence – is good for, or at least not bad for, the pregnant woman, her partner, her already-existing and her future offspring, her parents, her grandmother and *each other person who ever does or will exist*. Is it plausible to think that in that sort of case the early abortion is wrong? And if we think that it is wrong – in virtue of what it does to the merely possible person, the person whose coming into existence the early abortion will prevent – then don't we have to say the same thing about non-conception, and about conception followed by non-fruition? Doesn't (3) take us in the direction of an *implausibly stringent procreation obligation*?

Questions like these show that (3) is at least controversial. At the same time, however, the argument from (1) and (2) to (3) may strike us as perfectly valid. Surely, that is, moral law will insist that early abortion – and non-conception, and conception and non-fruition – is often wrong, if it has already been established both that merely possible persons matter morally in just the way you and I do, and that the choice of early abortion imposes a very great and deep loss on one person while saving still others from losses that are, in any ordinary case, relatively minor. Given (1) and (2), can there really be any moral distinction to be drawn between ending the life of an already-existing person and not bringing a new person into existence to begin with?

I want to suggest, however, that the argument is *not* in fact valid. We can accept (1) and (2) and reject (3). We can recognize, that is, that the merely possible matter morally in just the way that you and I do. And, for any case in which the existence itself is worth having, we can accept that not bringing the merely possible person into that existence is to impose a *loss* on that person. But we can at the same time maintain that early abortion – and non-conception, and conception followed by non-fruition – is perfectly permissible in many cases and indeed in practically all the cases when it is what the woman herself wants.

1.7.3 We will return to the validity of the argument below. It may seem that the simpler strategy for avoiding (3) would be to argue that (1) or (2) or both are false. It may seem tempting, for example, to argue that the merely possible do *not* matter morally in *anything like* the way that you and I matter morally, and hence that (1) is false.

That is the approach that Exclusion takes. Exclusion may seem commonsensical. How can a person who *never exists at all* matter morally in *anything like* the way that you or I, or our already-existing or future offspring, or the needy children on the next block or the other side of the planet, or future generations who may find themselves the victims of global warming, matter morally? Since the merely possible do not and never will exist, why should we concern ourselves with their plights, especially in light of the fact that our doing so will often come at a certain cost to persons who do or will exist – persons, that is, who clearly do matter morally?

In the end, however, we are going to be forced to reject Exclusion. Premise (1), understood in a certain way, is going to have to be accepted as an elemental principle, a veritable pillar, of population ethics. To see this – briefly here, and in more detail in Chapter 2 – just consider what happens when we *reject* (1).

Suppose that we say that the merely possible *don't* matter morally, that the losses they incur have no moral significance whatsoever and that, accordingly, their interests are to be *excluded* from our calculations of whether a given act is morally permissible or not. Serious conceptual difficulties immediately arise.

First, we want to be able to say that it *would be wrong* to bring merely possible persons into existence and then, for no reason that is good enough, to torture them. Moreover, the fact that we know that choice is wrong, even if unperformed, tells us, as we deliberate prior to performance, that we ought stay clear of that particular course of action – that we ought not bring the merely possible into existence and then torture them even if doing so would make things somewhat better for some people who do or will exist and worse for none. But the view that merely possible persons do not matter morally – that their losses are devoid of any moral significance – seems to preclude any such analysis. It seems instead – when put together with the otherwise plausible permissibility theory – to lead us inexorably to the result that – for example – the "genocidal adventures of nonactual dictators" are perfectly permissible so long as their victims remain nonactual – that is, merely possible – as well.[26]

Now, one might try to write this result off – though not, I think, very effectively – on the ground that we need not be concerned when a view generates troubling permissibility results in the case of acts that are *never* in fact performed. But there is a second result that is just as troubling and cannot be dismissed in that way. The position that the losses incurred by the merely possible are to be excluded from our calculations of the permissibility of a given act means, not just that our evaluations of our unperformed (nonactual) acts go awry, but also that our evaluations of our performed (actual) acts go awry as well. We will then be saddled with the result that what we have in fact done is wrong when all we have done is – for example – make things somewhat worse for a few people who do or will exist in order to avoid bringing distinct, merely possible persons into existence and then torturing them.

These two results indicate that it is *far* more difficult than one might have supposed simply to deny premise (1).

---

[26] This example is from Caspar Hare. See Hare (2007), p. 503.

1.7.4 Premise (2) as well seems a tempting target. And (2) is certainly controversial. For it supposes that we can assign a loss to a person at a world where that person never exists at all, which in turn seems to require that we can cogently say that an existence worth having is better for that person than never existing at all.

But (2) on closer inspection is hard to refute. To appreciate the case for (2), it is useful to have a bigger picture in mind. For purposes here, *losses* and *wrongs* are two quite distinct sorts of things. Thus, to say that an act (including any omission) imposes a *loss* on a person at a world is just to make a certain *metaphysical* point: that there is some other world accessible to agents such that that same person has additional wellbeing. On this view, we frequently incur losses – but whether those losses ground any particular moral assessment – that, for example, we have been wronged, or that a wrong has been done – is a different matter entirely. If wellbeing is taken away from one person who is already badly off just so that a distinct, already better off person can have still more wellbeing, the loss incurred by the one person may well signify that a wrong has been done. But in many other cases one or the other of two people will be forced to incur a loss yet we will say that no wrong has been done. Why not? Because we think that those losses have been traded off, or distributed between those two people, in a way that moral law itself directs.

Suppose, then, that we agree that an act imposes a loss on a person at a world just when there is some other world accessible to agents such that that same person has additional wellbeing. It is at least plausible that a person accrues no wellbeing at all, positive or negative – no benefits, no burdens, no goods, no bads, no plusses, no minuses – at any world at which that person never exists at all.[27] But at the same time it is clear that in many cases agents *do* have the alternative of bringing a person they leave out of existence at one world into an existence worth having at another world. From these two points, we seem easily able to infer that the act imposes, at the one world, a *loss* on the person it leaves out of existence.

There is nothing in this account that even begins to imply that the person who incurs the loss *exists* in both worlds. A person's existing in each of two worlds might be required in connection with some sorts of comparisons – when, for example, we say that a person is increasing his running speed faster at one world than another. But there is no reason to think that all comparisons will require the person to exist in both worlds. My mother, for example, need not exist in Beijing for it to make sense to say she has no money there – or for it to make sense to say that she has more money in Fort Worth, where she does exist, than she has in Beijing. Similarly, it makes sense to say that a person has more wellbeing at a world where that person has an existence worth having than at a world where that person never exists at all. If to incur a loss is just to have less wellbeing at one world than one has at another as a consequence of what agents have done at each of those worlds, then we have a loss.

We can even explain why the comparison itself is cogent, even though the person at issue exists at one world but not the other. What we are comparing is a *number* that

[27] This point is, of course, controversial. For further discussion, see part 2.2.2 below (the loss of never existing) and Roberts (2003a).

we have associated with one world (or place) and one person against a *number* we have associated with another world (or place) and that same person. On the basis of that comparison between *numbers* – all of which, unlike all *people*, exist in all worlds – we say that a loss has been imposed on that person at the one world.

One final point needs to be made about premise (2). (2) itself is at least *plausible enough* that we cannot comfortably let our acceptance of the moral permissibility of early abortion depend on the idea that (2) is false. *If* we think (2) is false, we are speculating. *If* (2) is false, it is false in virtue of some highly technical points having to do with the sorts of references that we can successfully make into worlds where the persons we aim to refer to never exist at all. Such issues would perhaps be worrisome in this context *except* for the fact that the merely possible persons we aim to be referring to clearly do exist in still other worlds. That means that any references into the worlds where those same persons never exist at all can be construed as references into worlds where they do exist. Even then, we obviously often won't be able to assign to those persons – who, say, relative to *this* world never exist at all – *genuine proper names*. Our references to them and our beliefs about them will not be *de re*. But that's often the case when a person has been or is about to be harmed. When shots are fired into a crowd and someone or another as a result is killed, we can judge that a harm, or loss, has been imposed on a person without ever having a clue just *who* has been harmed or incurred the loss. This is what quantifiers are for. Quantification and its related operations will let us say what we need to say.

1.7.5 In the end, then, (1) and (2) both seem perfectly correct. If so, then our only recourse for rejecting (3) is to reject the idea that (1) and (2) together compel us to accept (3).

And that, I believe, is exactly what we can and should do. In this connection, Variabilism is extremely useful. It shows just how we can accept both (1) and (2) but reject (3).

Let's start with (1). According to Variabilism, we each – you, me and the merely possible – matter morally, but we each matter variably. We can have things this way – we can find this middle ground – if we abandon the attempt to say who matters morally and focus instead on just which losses matter morally.

In other words, if our question is: "Does this merely possible *person* matter morally?" we have to say either yes or no, and both answers – the one that includes the merely possible person in the realm of those who matter morally and the one that excludes that person from the realm of those who matter morally – are problematic. If our question instead is: "Which among this merely possible person's *losses* matter morally?" we can say that *some* of this person's losses matter morally and *some* of this *same* person's losses do *not*. And we can say exactly the same thing about you, me and anyone else who does or will exist: some of our losses matter morally and some do not. Thus: the merely possible matter morally just as you and I do in respect of some of the losses they incur, but matter not at all, just as you and I matter not at all, in respect of still others.

What, according to Variabilism, imbues some of a particular person's losses with moral significance but not others – what makes some losses such that we *must* take them into account and other losses such that we must *not* take them into account in

calculating the permissibility of a particular act and its alternatives – is just where those losses are incurred in relation to the person who incurs them. Incurred in a circumstance, or possible future or world, where a person does or will exist, a loss will have full moral significance. Incurred anywhere else, a loss will have no moral significance whatsoever, not even the littlest bit, for purposes of evaluating the act that imposes that loss or its alternatives.

It's true that Variabilism humbles some among us in a certain way. If we ever vaguely thought that our own parents or other agents owed it to us to take whatever steps they could to bring us into existence to begin with – that just as they were required to create more wellbeing for us by not dropping us on our heads as infants, they were required to create more wellbeing for us by bringing us into existence – then we shall have to forego that claim. After all, though actual we undoubtedly are, had we never existed at all, the loss we then would have incurred by never having been brought into existence to begin with would have, according to Variabilism, *no moral significance whatsoever*. It would have been perfectly fine, in other words, for us never to have been conceived, born and (one hopes) adored at all.

This is to accept the idea that our coming into existence, while a matter of fabulously good luck, did not, in itself, make things better from a moral point of view. It has been, we can hope, "a blessing" for our parents, but it was never "a commandment."[28] But so what? The procreative entitlement we might vaguely have imagined to be ours in virtue of our status as actual persons against our parents is surely one we can live without.

Variabilism, thus, fully endorses premise (1): merely possible persons matter morally, just as you and I matter morally. We all matter variably – with the losses incurred by any person, merely possible or not, at worlds where that person does or will exist having full moral significance, and losses incurred at worlds where that person never exists having no moral significance at all, for purposes of evaluating both the acts that impose those losses and their alternatives.

1.7.6 Variabilism is, moreover, perfectly consistent with premise (2). In fact, Variabilism presupposes (2) – or at least is a far more plausible theory if (2) is true than if it is not. If we cannot recognize, on the basis of a comparison between the wellbeing one accrues if one never exists and the wellbeing one accrues if one has an existence worth having, that the early abortion imposes a loss on the person it causes never to exist, it is hard to see how we can recognize a loss, on the basis of an analogous comparison, in the context of wrongful life. If leaving someone out of an existence worth having does not impose a loss, neither does, it seems, bringing someone into an existence that is less than worth having. We should thus want to say, both for the case of early abortion – and non-conception, and conception followed by non-fruition – and for the case of wrongful life that agents have imposed a loss on a person in virtue of the fact that they have done less for that person when they could have done more.

1.7.7 And now we come to (3), the conclusion that early abortion is often – or ordinarily – wrong. Not all losses are, according to Variabilism, morally significant.

[28] Heyd (2009), p. 16.

Some have moral significance, and some do not. As noted, what determines whether a loss has moral significance, according to Variabilism, is just where that loss is incurred in relation to the person who incurs it. Incurred at a world where a person does or will exist, a loss will have full moral significance. Incurred anywhere else, a loss will have no moral significance whatsoever.

We thus rightfully insist – according to Variabilism – that the loss incurred in the context of wrongful life has full moral significance. For in that case the person who incurs the loss – the person born, for example, with the serious and unavoidable genetic or chromosomal disease – does or will exist at the same world at which the loss is incurred. But we can also insist – again citing Variabilism – that the loss incurred by virtue of the fact that a particular person never exists at all is devoid of moral significance, since that loss is incurred by that person at a world where that person never exists at all.

And thus we can accept (1) and (2) and reject (3). We can say that the person whose coming into existence is thwarted does indeed incur a loss. But we can also say that that loss has no moral significance. And the loss – under any otherwise plausible permissibility theory – that has no moral significance at all cannot count against the act that imposes that loss. A loss that has no moral significance at all cannot make an otherwise permissible act wrong or, in a roundabout way, make an otherwise wrong alternative to that act right.

1.8 Tradeoffs and Abortion

1.8.1 So far, I have only briefly sketched the conclusions I aim to reach regarding early and late abortion – that early abortions are ordinarily permissible, and under very summary analyses, and late abortions are ordinarily very hard to justify. What I have said almost nothing about so far are the myriad *tradeoffs* that can arise in abortion scenarios. There will be more to be said about tradeoffs in Chapter 5 below. But it is important to say just a bit here to avoid creating any deep misconception starting out.

Let's first consider the early abortion – the choice not to bring a person into existence to begin with. Many cases will not be complicated by the fact that a *tradeoff* in favor of one party or the other will need to be made. In many cases, that is, no one will have any stake in how the choice of early abortion is made *other than* the person whose coming into an existence worth having is at issue and, perhaps, the pregnant woman herself. In such cases, Variabilism, in combination with the otherwise plausible permissibility theory, immediately implies that the early abortion is perfectly permissible. The early abortion is permissible because it avoids morally significant losses on behalf of, or at least does not impose any morally significant losses on, *anyone at all*. It's true that the person left out of existence altogether as a result of the early abortion incurs a loss that is very great and very deep. But, according to Variabilism, that particular loss is devoid of moral significance. The otherwise plausible permissibility theory, accordingly, will not take that

loss into account in evaluating the early abortion – which will in turn clear the way for the result that the early abortion is itself permissible.

In any such case, the pregnant woman may have the early abortion for a good reason, for a poor reason or for no reason at all. She need not prove or believe that her fetus is defective, or argue that her health is at stake, or claim or provide evidence that she is the victim of rape or incest. She may do just as she pleases, provided that any abortion takes place within that span of time nature allows between the moment of conception and the moment at which the organism growing within her uterus must be recognized as a thinking and otherwise continuing thing.

1.8.2 But in still other cases the losses that the early abortion stands to impose will have full moral significance, according to Variabilism. The woman's partner, for example, and even her *grandmother* may incur just such losses if the early abortion is in fact performed, since those losses will be incurred at worlds where the persons who incur those losses – the partner; the grandmother – do or will exist. In such cases, we do face *tradeoff scenarios*. And we face *morally significant* tradeoff scenarios at that. For the tradeoffs are not just between one loss that has full moral significance and another that has no moral significance at all, but rather between two losses that each have full moral significance.

Such tradeoffs are never completely straightforward. What we are to do about them is a matter that neither Variabilism nor any of its competitors can help us with. They are matters to be addressed, instead, within the scope of the otherwise plausible permissibility theory.

Nonetheless, we can predict that the otherwise plausible permissibility theory in many cases will imply that the right way of making the tradeoff between the woman and her partner, or the woman and her grandmother, will be for the woman to have the early abortion. For one thing, very often the need to make the tradeoff will be illusory. There will be a way to avoid imposing a morally significant loss on anyone at all. There will be a way to accommodate the interests of the woman's partner or her grandmother, short of the woman's continuing the pregnancy she does not want. Moreover, when the conflict is genuine – that is, when the imposition of some morally significant loss or another is *unavoidable* – the pregnant woman will typically face, if not the demonstrably greater loss, then at least the demonstrably greater *risk* of loss. And risk may well be what is pertinent for purposes of assessing permissibility prior to choice, since we will often have scant information regarding actual future loss for any of the parties involved. Here, too, the otherwise plausible permissibility theory may well instruct that the tradeoff is to be made in favor of the pregnant woman.

But in other cases we will reach different conclusions. Thus, we can imagine cases where the loss that the woman will incur is negligible, and the loss her partner, or her grandmother, faces far more serious. We can thus certainly imagine cases in which the otherwise plausible permissibility theory will instruct – as a matter of moral, not positive, law – that the woman ought not have the early abortion.

1.8.3 A different treatment of late abortion is in order. The pregnancy has been allowed to proceed to the point where the fetus has had that first thought that signals a person's coming into existence. And that person has acquired the capacity to survive from one moment to another as the same person. Accordingly, there is no argument

to be made that the late abortion will simply cause an existing person to lose a future that that person would have lost in any event. And let's suppose that we are not in a case where that future would be *less* than worth having. We are not, that is, in a case where the late abortion rescues a person from a genuinely wrongful life.

Accordingly, the late abortion will cause the newly existing person to incur a loss. And, since that loss is incurred at a world where that person exists, according to Variabilism, that loss itself must be recognized as having full moral significance.

Of course, the pregnant woman may well incur a morally significant loss if the late abortion is *not* performed. She, after all, exists, too – and her losses count as well, according to Variabilism.

Nonetheless, it seems that, in any ordinary case, the otherwise plausible permissibility theory will instruct that the tradeoff is to be made in favor of the newly existing person rather than the pregnant woman. The morally significant loss that the newly existing person will incur if the late abortion is performed is both very great and very deep. A whole lifetime's worth of wellbeing will have been lost – and it will be lost by a person who has had very little chance to accrue very much wellbeing to begin with. In contrast, the pregnant woman will often already have accrued a good deal of wellbeing, and she may well expect to accrue a good deal more whether she has the late abortion or not.

On these various moral data, it seems that an otherwise plausible permissibility theory will say that the tradeoff between the pregnant woman and the newly existing person must be made in favor of the new person. There will always be the philosopher's case in which the woman must have the late abortion to save the world – or the five human (or non-human; what is important is that they are *persons*) neonates who happen to be cradled alongside the trolley track. There will be the more realistic case in which the "pregnant woman" herself is a child whose own life to date may itself have been barely worth living (or even wrongful, that is, less than worth living) and whose own future may be very bleak. But in any more ordinary case, the otherwise plausible permissibility theory will deem the late abortion wrong.

1.9 Abortion and the Law

Until we have the otherwise plausible permissibility theory firmly in hand, it might seem that we really cannot know what that theory will say about the choice of late abortion in any case where so many other agents are doing things – lawfully, and apparently with the full blessing of society – that are at least as bad and getting away with them.

I suspect that an otherwise plausible permissibility theory will not consider the fact that most other agents are behaving very badly a justification for any one agent's behaving very badly. That everyone else is doing it may explain from a psychological point of view why I've done something wrong. But it seems that, under any otherwise plausible permissibility theory, it won't make the otherwise wrong act morally permissible.

But does it make it *legal*? It could be argued that the bare fact that persons other than pregnant women are legally permitted to do very nearly as they please to many human beings and virtually all non-human animals means that a law requiring pregnant women to adhere to a far more stringent standard violates constitutional equal protection guaranties.[29] Moreover, the fact that it is wrong for the woman to choose the late abortion cannot, on its own, imply that it is permissible for the majority to force the pregnant woman to do the right thing when they themselves in analogous cases are left to decide for themselves. But these are just arguments – and difficult arguments, at that. The better view is probably that the U.S. Constitution – for example – should be understood to protect the woman's right of early abortion but not late abortion. Read in this way, the Constitution would align the cases in just the way Variabilism suggests, recognizing a loss the state has a compelling interest in preventing in the case of late but not in the case of early abortion.

1.10 A Middle Ground on Abortion

The middle ground I propose on abortion thus takes a very hard line on late abortion and almost no line at all on early abortion. Why is *this* middle ground? Is what I am calling *middle ground* in fact going to offend practically *everyone*? I don't think it should. Surely, political conservatives and religious leaders who disdain abortion at any stage of pregnancy for any reason are most concerned to see abortion recognized as impermissible and to see it banned by law late in pregnancy. My proposal reflects that concern. It considers late abortion – not moments before birth but any time after the person – the thinking thing – has come into existence and acquired the capacity to survive – almost always wrong. Neither serious fetal abnormality (barring the case of the genuinely wrongful life), nor maternal health (physical or psychological), nor family circumstances (poverty or resources stretched thin by the needs of existing or future offspring), nor the conditions of conception (rape or incest), nor the woman's own ignorance of the fact that she is pregnant, nor the 16-year-old girl's status as a minor child, would be enough to justify the late abortion.

By the same token, surely those who think abortion is permissible throughout the pregnancy, including those who consider abortion a matter of bodily autonomy, are most concerned to have that fact recognized in connection with early abortion. My proposal reflects that concern as well. Early abortion is almost always perfectly permissible in view of the fact that the single substantial loss that we can expect the abortion to impose in any case in which it is what the woman wants – the loss incurred by the person whose coming into existence is prevented by the early abortion – will be deemed by Variabilism to have no moral significance whatsoever.

[29] At issue here would be, for example, the equal protection clause of the Fourteenth Amendment of the U.S. Constitution. See part 5.7 below.

My proposal thus effectively asks each of the two camps settle for the ground that is most important to it and concede the ground that is most important to the opposition. The discussion of early abortion will no doubt remain fraught – for political or psychological reasons, or because of an evolutionary instinct many men and no doubt some grandmothers have in common that strongly inclines them to find in moral law edicts commanding women to have babies. The discussion of late abortion may remain fraught as well. In recent years, moral philosophers, by virtue of the fact that they aren't quite clear what to say about the moral status of merely possible persons, have only contributed to the confusion. Their theories either weigh against even the earliest abortion or are so subtle we are not quite sure exactly what is being said at all. My goal here is to unwind things just a bit, by first coming to a clearer understanding of the moral significance of merely possible persons and only then beginning to address the problem of abortion.