

Supernumerary Pregnancy, Collective Harm, and Two Forms of the Nonidentity Problem

M. A. Roberts*

I.

A serious risk associated with some treatments for infertility is multiple birth pregnancy. The term “supernumerary pregnancy” aptly conveys the idea of *too much* or *too many* in this context: too many fetuses developing in the womb to secure a reasonable probability of a good outcome for any. A supernumerary pregnancy can occur, for example, when infertility is treated by the aggressive use of ovulation-stimulating hormones. If pregnancy is then achieved either naturally or artificially, or if large numbers of eggs are retrieved, fertilized in vitro, and then transferred to the uterus in a single cycle, the risk of supernumerary pregnancy is increased. Egg donation coupled with artificial insemination, or straight embryo donation, followed in each case by the transfer of more than two or three embryos to the uterus also increases the risk for supernumerary pregnancy.

An interesting question is whether the babies born of an infertility treatment-induced supernumerary pregnancy (ITISP) are properly considered to have been *harmed*. One might wonder how such a question could even arise in the face of data that clearly demonstrate that ITISP leaves an unduly large number of babies blind, deaf and palsied and facing life-long disabilities.¹ On the face of things, these babies seem to have been obviously badly damaged by

the very health care providers and parents charged with their care. In fact, however, a number of arguments seek to undermine any clear claim that harm has been done within the ITISP context.

The purpose of the present paper is to examine three such arguments. The first is rooted in longstanding legal debates regarding the conditions under which an individual agent can be said to have harmed a person when that person's plight is the work of the combined activities of multiple agents. The latter two derive from more contemporary discussions within ethics and the law regarding when, if ever, agents can be said to harm a person by making the procreative choice to bring that very same person into existence. The former emerges as the issue of how the phenomenon of *collective harm* is to be analyzed. The latter two appear within philosophy as distinct forms of the always beguiling *nonidentity problem* and within the law in connection with claims in negligence for *wrongful life*.

In opposition to these “no harm done” arguments, I will propose a criterion of harm that, at least at the root, fits together with both (1) a plausible (though not uncontroversial) approach within the law to the problem of collective harm and (2) a core insight (the “basic maximizing insight”) behind most contemporary consequentialist approaches to the question of how acts are to be assessed from the moral point of view. I will then argue that the proper application of that criterion indeed enables us to discern harm done to ITISP-damaged babies – a fact that I believe will have, and should have, significant moral and legal implications.

The paper is organized as follows: I briefly note, in Part II below, how the seemingly mundane question of harm pertains to critical and urgent issues relating to the structure of morality, the relation between law and moral theory, and the problem of indeterminacy in the law.

In Part III, I argue that it is a serious mistake to think that ITISP can be assimilated to the standard “wrongful life” model. A claim for wrongful life may be brought on behalf of an impaired child when a health care provider has negligently failed to advise the child’s parents that they face an increased risk (attributable, for example, to parental age or family history) of producing a child with a genetic or chromosomal abnormality. Most courts, however, have rejected claims for wrongful life on the grounds that, as a matter of law, it cannot be established that the child has been harmed or has suffered damages. After all, the child’s flawed existence is the child’s *only alternative* to the child’s never having been born at all; and courts are reluctant to admit that nonexistence can ever be better for the child than a flawed existence.² Since the only way to avoid the genetic or chromosomal abnormality is to produce some distinct child, “nonidentical” to the original, or no child at all, the wrongful life model in effect presents us with a first version of the nonidentity problem. I acknowledge that the standard wrongful life model plausibly applies to some applications of assisted reproductive technologies (ARTs), including applications recently associated with increased risks of offspring disability.³ I shall argue, however, that that same model applies to virtually *no* instance of ITISP.

Of course, whether harm is imposed in the ITISP context is in part a question of what account of harm we adopt. I thus will argue, also in Part III, that the account of harm we most plausibly put to work in the ITISP context incorporates the very criterion put to work in the standard wrongful life model: assuming the existence is one worth having, we find “no harm done” if, but only if, being brought into existence in an impaired state is the child’s *only alternative* to never having existed at all. In effect a *maximizing* test for harm, the idea, roughly speaking, is that when we see just how we *could have done more* for a given person than we have

and we in fact do less, it becomes plausible to say we have *harmed* that person. My argument will be that it is clear in the ITISP context that agents could, indeed, have done more for the damaged offspring than they actually have. It thus follows easily, I argue, that such offspring have in fact been harmed.

Part IV briefly examines a competing, narrower criterion of harm that often seems at play in discussions of procreative choice. According to that narrower criterion, to confer a worthwhile existence on a particular person is to confer a *benefit* on that person, and any effect that constitutes a *benefit* cannot also cogently be considered to be a *harm*. I argue, in Part IV, that the second part of this narrower criterion is problematic, and that it is critical to distinguish between what it is to do “some good” for a person and what it is not to harm that person. In fact, it is easy to see that we can do “some good” for a given person yet nonetheless harm that person – and thus that conferring the benefit of existence on a given person cannot, on its own, suffice to show “no harm done.”

More serious challenges against the preliminary account of harm I propose in Part III must, however, be met as well. That account fails to explain how the idea that the agents – the health care providers and the parents – *together* could have done more for the damaged babies than they have translates to answer questions about whether acts performed by the agents *as individuals* can be considered harmful. Does an individual agent impose harm on a given person *only if* that agent has the ability, on his or her own, independent of any choices made by any other agent, to have made things better for that person than they are? Where, for example, the health care provider has done his or her best for a given child but only in virtue of the fact that the provider’s options for doing still more have been effectively eliminated by the patient in the exercise of her autonomy in respect of her own pregnancy, has the *provider* harmed the child? It may seem that the provider,

who has done the best he or she could have done for the child, clearly has not *harmed* the child; it may seem that the harm that has been inflicted by the *group* consisting of provider and patient surely cannot be properly attributed to the *provider*. Cases involving *reciprocally limited options*, however, take us just as clearly in the direction of the view that harm that only the group can avoid may after all be attributed to the agent in certain cases. Just as the patient may limit the *provider's* options for doing more for the child, so may the provider effectively limit the *patient's*. Can we really plausibly say, about such a case, that *no person*, only the group, has harmed the child?

Such questions are the subject of a longstanding debate within the law of torts, and the goal of Part V is to propose a way of narrowing that debate. I thus suggest, in Part V, a criterion of harm that retains a maximizing element but that also has the potential to attribute group-harm to individuals in some cases.

Part VI addresses a second form of the nonidentity problem – a powerful, probabilistic form of that problem that can be generated in the ITISP context. Suppose that health care providers and parents had avoided ITISP and instead done things in the supposedly “better-for-the-baby” way I prescribe in Part III. It might then seem that the odds of that very baby having ever come into existence at all would have been quite long. Surely, my prescription would have led, not to the one baby’s being made better off, but rather to a better-off baby, a baby “nonidentical” to the original, having been made.⁴ But it is surely better to secure a flawed existence than to be accorded a very long shot at an unflawed existence. It may therefore seem that, from the baby’s own point of view, ITISP is not a bad thing and does not harm that baby at all. The purpose of Part VI, then, is to disarm this second form of the nonidentity problem. I there argue that the defect in the argument is a matter of the relevant probabilities having been misstated. Once that

misstatement is identified and corrected, this second form of the nonidentity problem, I argue, collapses as well, allowing us to discern the harm that has in fact been done to the baby.

Conclusions are drawn in Part VII.

II. Harm, Law and Moral Theory

Part of the reason I focus on the question of harm in the context of ITISP in this paper and in other contexts elsewhere is that I am interested in the connections between law and moral theory and, in particular, the problem of indeterminacy in the law. Suppose we think that it is *morally wrong* for agents to work together to produce ITISP-damaged offspring but also think that no single member of the group of agents who have *together* engineered an instance of ITISP has harmed any offspring then produced. Perhaps we have come to that assessment because we think that harm agents together impose cannot properly be attributed to the agents as individuals. Perhaps, alternatively, we have come to that assessment because we think that acts that bring the very people whom they make suffer into a worthwhile existence cannot harm those people. Either way, it is our assessment that ITISP proves that *moral wrongdoing* can take place even in the absence of any agent having *harmed* any person. We shall then be forced to reject the so-called “person-affecting,” or “person-based,” intuition – the intuition, that is, that “what is bad must be bad *for someone*” and that wrongdoing can occur only on the condition that someone, somewhere, have been *harmed*, or *made worse off*, by what the agent has done.⁵ Now, often it is a good and clarifying thing to reject some “intuition” one has previously clung to, as where that “intuition” constitutes an unexamined feeling, or bias, one has accepted as gospel but that in reality is in some way obfuscating. But there is reason to think that rejecting the person-based intuition falls into

another category. Rejecting that particular intuition may not clarify things at all but rather create a great divide between law and morality, a divide we may in the end find very hard to live with.

Within the very branches of the law that govern procreation – tort law, constitutional privacy law and family law – the person-based intuition is deeply entrenched. Attempts to use processes of law to regulate or prohibit conduct will thus always be misguided, in the eyes of the law, when agents – in some cases the defendants in a negligence action and in other cases individuals seeking to exercise their presumed constitutional right to live the lives they consider to have value – can be definitively shown to have imposed *no harm*, and *no risk of harm*, on *any* existing or future person.⁶ In this context, our adoption of an “impersonal” moral theory – a moral theory that rejects the person-based intuition and instead provides that wrongdoing has no substantial connection to the harming of persons – will mean that the results we are able to derive from that moral theory may have *no relevance whatsoever* within the law. What morality deems “wrong,” the law of negligence – which requires its causes of action to contain references to duty, breach, cause *and harm* – will consider the “failure to state a claim.” What morality deems “wrong,” the Fourteenth Amendment – which requires states to give individuals the freedom to make certain intimate and important choices for themselves – may consider the exercise of a constitutional right of privacy. A deeply personal choice of critical importance to the individual that demonstrably imposes *no harm* and *no risk of harm* on any existing or future person? Whatever morality might say about such a choice, a court may be reluctant to find any compelling reason for the state’s prohibition of it. The ships of law and moral theory will pass in the night.

This is not a divide we should love. Rather, it should alarm us. The law itself remains highly indeterminate with respect to many hard questions regarding future persons – including,

among many others, maternal-fetal conflicts, the new genetic and reproductive technologies, ownership and custody of frozen embryos, human reproductive cloning, and abortion. In the face of that very deep indeterminacy, addressing those questions in anything other than an arbitrary way *without* bringing moral principle to bear is a daunting task.

Of course, if we think that moral law truly is impersonal in nature, we always have the option of restructuring the law along impersonal lines.⁷ If harm is not an essential element of a moral wrong, then we can close the divide by revising the law – that human artifact – so that harm becomes inessential as well for purposes of articulating a legal wrong or a constitutional basis for state regulation. We can then say that, just as ITISP constitutes a moral wrong, it also constitutes a legal wrong – a case of medical malpractice, for example, or a case of parental abuse or neglect, or a sound constitutional basis for state or federal regulations aimed at reducing the risk of ITISP – despite the fact that ITISP actually *harms* no one.

Such a complete restructuring of the law, however, is a dramatic step. My own hesitation in taking this step is that I cannot see how we can be so sure, at this early stage of the investigation, that it is the law rather than moral theory that has made the mistake. For one thing, moral theories that take an impersonal form – theories, for example, that determine wrongdoing on the basis of whether the agent has maximized *aggregate*, or *total*, wellbeing – face their own very grave challenges.⁸ Moreover, as I argue in the parts that follow, it may well be a mistake to think that we cannot successfully discern harm in the context of ITISP. If I am correct, then, in effect, we will have no clear need to move from an intuitive, person-based approach to morality to a highly problematic impersonal approach.

A second strategy for closing the divide between law and moral theory would be to redefine the term “harm” in a way that would function to increase the prevalence of harm in the world. Thus, we might define “harm” so that ITISP imposes harm in some stipulative sense even if it fails to impose harm in any ordinary or intuitive sense. Then, the mere fact that a baby is born blind or with any other serious impairment or that the baby’s quality of life falls below a certain threshold would entail that bringing the baby into existence “harms” that baby, even in the case where the only alternative for that same baby is never having existed at all.⁹ Alternatively, we could define “harm” so that agents who had the option of creating a better off baby in place of an impaired baby would automatically be deemed to have “harmed” the latter.¹⁰

For many reasons, however, this second strategy is problematic as well. Most importantly, accepting a stipulative definition of “harm” runs the risk of disorienting our moral compass. The “harm” of a badly-mangled eye surgery that inexcusably leaves the patient blind and the “harm” of having a genetic predisposition for blindness run together. We are left hard pressed to say why it is that one is actionable and the other merely unfortunate. Moreover, while this second strategy on its face closes the divide between moral and legal theory – at least, that is, when the law demands that *harm* be shown, the moral theorist will be in a position to show “harm” – it is highly doubtful that it is a strategy courts will generally accept. Practically speaking, it seems more likely that courts that do not see how ITISP-damaged babies have been harmed in an intuitive and ordinary sense will be unwilling to take the position that the existing tort and family law frameworks can be put to work to help protect the interests of those babies, or that, consistent with the Fourteenth Amendment, states may act to protect future offspring against such damage.

Still a third strategy for closing the divide between law and moral theory is simply to accept that, despite strong sentiment to the contrary, there truly is nothing morally wrong with agents – parents and health care providers – working together to bring into existence, via ITISP, up to, say, eight badly damaged babies in a single pregnancy. But this third strategy may well be considered the least appealing among our available options. If morality deems *that* choice permissible, then it is hard to imagine what sorts of procreative choices morality will rule out. More plausibly, a correct moral theory will be one that sets fairly stringent limits in the context of procreative choice.

Compared to each of these three strategies, the project of showing that ITISP does after all involve genuine (non-stipulative) harm to offspring is a project we may well all hope will succeed. The question, then, is whether such a showing, in the face of issues relating to collective harm and the nonidentity problem, can cogently be made. Has the baby left blind, deaf, palsied and disabled for life by way of ITISP been *harmed* in some intuitive, ordinary sense of the term “harm”?

III. The “Wrongful Life” Version of the Nonidentity Problem; Why ITISP is Nothing Like Wrongful Life

When faced with claims for wrongful life, most courts have taken the view that the procreative effect of the health care provider’s negligence – life itself – does not constitute a harm to the child. In the absence of any finding of harm, most courts, correspondingly, have rejected the child’s claim for damages.¹¹

This way of looking at wrongful life seems roughly plausible. Much of the child’s suffering in the wrongful life scenario is due to the child’s being born with a chromosomal or

genetic defect. Nothing the provider or the parent, as individuals or as a group, could have done at any time would have been enough to keep that particular child from having been born with that particular genetic or chromosomal defect. Agents could have “saved” the child from suffering, to be sure, but only by not bringing that child into existence at all and instead bringing a second, healthier child – a second child *nonidentical* to the child who suffers – or no child at all into existence in place of the one. On the assumption that the particular case is one in which the flawed existence is preferable to none at all, the child is plausibly is not harmed when that flawed existence is negligently conferred.

But ITISP is nothing like wrongful life. In ITISP, the damage to offspring is a function not of a genetic or chromosomal defect but of uterine overcrowding, engineered not by nature but by health care providers and their patients. And what has been engineered *in* often can – and could have been in this case – engineered *out*. Thus, the provider could have put safeguards against supernumerary pregnancy into place – safeguards that would have enabled the *very same babies* in fact born impaired (or some of them in any event) to have enjoyed a *better* existence. Using egg retrieval and IVF, the provider thus could have controlled either (1) the number of embryos created or (2) the number of embryos transferred to the uterus. Alternatively, (3) once a supernumerary pregnancy is confirmed by sonogram, fetal reduction, in which the fetal heart is stopped by an injection of potassium chloride in the tenth to twelfth week of the pregnancy, can be recommended and made available to the patient. In each case, of course, the patient has the option of going along with this advice.

It is clear, then, that, for each baby who suffers from an ITISP-related impairment, the provider, working together with the patient, had the option of saving that baby from impairment.

Provider and patient could have given that same baby a better existence by opting to bring fewer babies into existence to begin with. Of course, the patient and provider could not have created additional wellbeing for *all* the babies born of the supernumerary pregnancy. But they did have the option of rescuing any one, and perhaps the option of rescuing any two or three. When provider and patient then fail to take that option, it becomes plausible to say the impaired child has indeed been harmed by the choices that were made.¹²

In distinguishing wrongful life from ITISP, it is critical that the various safeguards against supernumerary pregnancy just mentioned are “identity-preserving” in nature. The safeguards, in other words, represent options for bringing *the same, identical child* into existence but free of the disability. Thus, in wrongful life, what nature has engineered into a given child cannot be eliminated without leaving that particular child out of existence altogether. In ITISP, in contrast, agents have the option of bringing the same, identical child into existence in an unimpaired state. Correspondingly, in wrongful life, a child exists for whom the provider and parent, working together, have done as much they can. In ITISP, in contrast, a child exists for whom the provider and parent, working together, could have done much more than they have.

The preceding discussion focuses on the options agents have *together* for insuring the wellbeing of a particular child. The fact, however, that the agents *as a group* have a particular option does not imply that each individual member of that group has that same option. That agents *as a group* – provider and patient working together – could have put one or more identity-preserving options to work to increase wellbeing for a particular child does not imply that each agent operating as an *individual* could have likewise acted to increase wellbeing for a particular child. Suppose, for example, that the provider fails to inform the patient of the risk of ITISP or

fails to explain the various safeguards for addressing that risk in a timely way. Suppose, as well, that had the patient understood the risk, she would have gladly consented to one of the options for addressing that risk. We might then say that the provider *as individual* had the option of creating additional wellbeing for the baby while the patient *as individual* did not. Similarly, if the provider does an excellent job conveying the risks and safeguards to the patient and the patient refuses to put any of those safeguards to work, then we might say that the patient *as individual* had the option of creating additional wellbeing for the baby while the provider *as individual* did not.¹³

The following criterion of harm allows us to make the critical and intuitive distinction between wrongful life and ITISP, taking into account that agents as a group may be able to achieve what agents as individuals cannot. Thus:

Maximizing Criterion of Harm (First Version) (MCHI).

Where p is any existing person, harm is imposed on p :

1. by an *individual* agent i , if i had the option of creating more wellbeing for p and instead created less; and
2. by a *group* of agents $i_1 \dots i_n$, if $i_1 \dots i_n$ together had the option of creating more wellbeing for p and instead created less.

Such a person p may already be in existence at the time of action or may at that time be a future person, that is, someone who will, but has not yet, come into existence. In the latter case, harm is imposed only if and only when p comes into existence. “Acts” include affirmative acts as well as omissions. The term “wellbeing” is left undefined, with the understanding that wellbeing corresponds in some rough way to whatever it is that makes life precious to the one who lives.¹⁴

Suppose, then, that in a particular case the upshot of agents opting for ITISP is a group of eight badly damaged babies born fifteen weeks premature and weighing between nine and twenty-six ounces – a group of babies who, if they survive at all, will face a lifetime of physical and mental

disability.¹⁵ It is, at the very least, clear that the provider and patient acting as a group could have done more for any one of the eight babies than they have. This is so, since, as we have seen, the provider and patient together have the option of putting identity-preserving safeguards against ITISP into place. Clause (2) of MCH1 thus implies that a given baby has been harmed by the group consisting of both provider and patient. It may be true as well that the provider or the patient acting as individual had the option of putting identity-preserving safeguards against ITISP into place. If so, clause (1) of MCH1 will imply, in addition, that the baby has been harmed by that individual.

The same inference holds for each of the remaining eight damaged babies born. For each such baby, in other words, agents, either as a group or as an individual or both, had options for creating additional wellbeing for that baby. And that is enough to establish, under MCH1, for each and every baby among that group who exists and suffers, that harm has been done.

Two points require clarification. First, consistent with my account of ITISP, just how the provider and the patient, working together, go about not harming the babies remains a matter of choice for those agents. That the option of putting one or more identity-preserving safeguards – egg retrieval, limited insemination, fetal reduction – into place exists is enough to prove, under MCH1, that harm has been done. But that such an option *exists* does not imply that agents are required to make *use* of it. This is so, in virtue of the fact that the agents, plausibly, have no categorical obligation to make things better for the very babies whom they, according to MCH1, harm, since they, plausibly, have no obligation to bring *those very babies* into existence at all. Since agents have no obligation to make things better for *those very babies*, agents have no obligation to use an *identity-preserving* option in order to make things better for them. Instead, the

agents' obligation is to avoid imposing harm on whatever babies – whether the babies who do in fact exist and suffer, or some other collection of babies altogether – they happen to bring into existence. And that obligation they can discharge by using any *non-identity-preserving option* for reducing the risk of supernumerary pregnancy that they please. Such options include, for example, (1) reducing the dosage of the ovulation-stimulating hormone and (2) using sonogram to establish whether an excessive number of follicles are maturing in a given cycle, and if there are, delaying intercourse (or artificial insemination) until the next cycle. Other options exist as well – including the option of avoiding the pregnancy altogether. Each such option, of course, would have dramatically reduced the chances of existence for each baby who in fact exists and suffers ITISP-related damage. But from that fact it does not follow, under MCH1, that the use of such a non-identity-preserving option would have harmed *that* baby – or any other. MCH1, thus, is expressly limited in its application to *existing* persons. The theorist who wants to show that we harm those merely *possible* people whom we never bring into existence will require another principle. Is such a supplementary principle plausible? I do not think so. We too clearly distinguish between not bringing a person into existence to begin with and removing a particular person from existence. If we harm in any sense at all those who do not and never will exist, the harm that we do them is surely not morally, or legally, significant.

These points make the moral analysis a simpler matter. The identity-preserving safeguards against ITISP tend to be expensive, and some of them may be offensive to the patient's moral sensibilities or run counter to her religious commitments or both. When expense is an issue, or when moral or religious sentiment is at odds with an identity-preserving safeguard, our discussion shows that agents are quite free to opt for a non-identity preserving safeguard instead. Thus, harm

to person *p* is established on the basis of whether better options for *p* exist. And better options – identity-preserving options – do exist in this case. But the question how agents want to go about avoiding that harm is a completely independent question. Consistent with MCH1, then, an option for avoiding harm to *p* is to put a non-identity preserving safeguard to work even though that option may well mean that *p* never comes into existence at all.

Second, it might be objected that MCH1 – and a maximizing approach to harm generally – is surely overbroad. For example, I harm Bill Gates, under any such approach, when I fail to send him the contents of my savings account; I harm the beneficiaries of Save The Children when I fail to double the amount I sent that organization the month before.¹⁶

In fact, however, any concern that such an account of harm is overbroad should evaporate under close inspection. After all, under the law, harm is just one of several necessary elements of a claim of negligence. The plaintiff must also establish duty, breach and causation and be able to establish the amount of damages the defendant is claimed to owe in some reasonably precise way. Moreover, any moral theory that assesses acts by reference to the consequences they have *for people* – a category that includes many deontological theories as well as virtually all contemporary forms of consequentialism – will consider the question of harm likewise just to be one component along the path to determining whether an act is wrong. Where an act that harms one person also creates more aggregate wellbeing than its alternatives (on an aggregative, or total, form of consequentialism), or is necessary to avoid imposing a still worse harm on someone else (on a person-based form of consequentialism), then that act may well be morally permissible – that is, not wrong – despite the harm that it does. In the domains of both law and ethics, then, *harm* is a matter that is closely intertwined with *causation*: it involves the metaphysics of making someone

worse off than he or she might otherwise have been. We do not cross the line between the metaphysical and the moral realms until we ask further whether what the agent has done is *wrong*.

Distinguishing questions of harming and wronging, then, we can almost automatically and instantaneously compute that, while I may have harmed Bill Gates, I certainly have not wronged him. Similarly, when I harm one child by failing to enroll that child in an expensive private nursery school so that I might feed my other children, we can with just a little thought determine that the harm I have imposed on the one child does not amount to a wrong in view of the fact that it was necessary to avoid the imposition of still graver harms on others. (Here, the totalist might point out that imposing harm on the one child was necessary to avoid a diminution in *aggregate* wellbeing. The person-based account would be a bit different. The analysis there would be that imposing harm on the one child was necessary to avoid a still greater diminution in *individual* wellbeing for another.)¹⁷

The distinction between harming and wronging gives rise to a further question more directly related to ITISP. A critic might concede that, as MCH1 suggests, each of the eight damaged babies is harmed by ITISP. But cannot the critic now argue that, though *harmed*, none of the eight babies have been *wronged*? After all, surely as between the flawed existence and no existence at all, the latter is the graver harm. And, as just noted, where imposing a harm on one person is necessary to avoid imposing still graver harms on others, the one harm may not constitute a wrong.

The problem with this argument is that it assumes that not bringing a person into existence to begin with constitutes some form of significant harm to that person. As noted before, however, and as MCH1 suggests, that view seems implausible. Rather, ITISP seems simply not to be the

kind of case in which agents are required to make a difficult tradeoff between imposing serious harm on one or two persons and imposing still graver harms on still more persons. ITISP, instead, is better understood to involve no serious tradeoff at all.

IV. The “All ARTs are Harmless” View: Another Way of Looking at ITISP

There is another way of looking at ITISP according to which ITISP is very much to be understood on the wrongful life model. Treatments for infertility are procreative in nature: when they work, they bring a new person into existence. As long as that existence is worth living, isn't such a treatment of necessity a *good* thing from the perspective of the one who lives? But if this is so, how can it also constitute a *harm* to that person? This line of analysis imagines wrongful life, ITISP and, indeed, all ARTs to be of a piece and suggests that the “no harm done” result always holds except perhaps in the case where the life is so full of pain and misery as to be less than one worth living.

Does anyone really think that an outcome cannot, on a net basis, constitute a good for a given person, yet still involve harm to that person? Does anyone really think that rescuing the heart attack victim by way of a heart transplant when a simple aspirin would have done the trick really does not harm that person even in the case where the heart attack victim is left with a life worth living, itself a good thing, and would have died in the absence of any treatment at all?

Surely not. More plausibly, passages that suggest otherwise have been drafted under the critical if not consistently explicitly articulated assumption that the ART under evaluation constitutes *the only route* into existence for a particular person. Thus, in the absence of that assumption, the following passage from Carl Coleman seems implausible:

Once the perspective of a particular child is assumed, it becomes virtually impossible to say that ARTs have "harmed" that child, even if the child is born severely impaired. Withholding ARTs would not have led to the birth of the child without the impairments; instead, if ARTs were withheld, the child would never have existed at all. The only situation in which it would make sense to say that a particular child has been harmed by being born with significant impairments is when the burdens of life with the impairments are so severe that existence itself is a net disadvantage – that is, if the child, once alive, would prefer that she had never been born.¹⁸

Now, the assumption that Coleman does explicitly make – that the case is one in which, “if ARTs *were* withheld, the child *would* never have existed at all” – is not adequate to ground the “no harm done” result with which he concludes. Thus, consider the statement that “if the surgeon doesn’t perform open-heart surgery, the heart patient will die.” This statement, let’s suppose, is true in virtue of the fact that the surgeon, if barred from performing open-heart surgery, would refuse, in a pique, to give any emergency treatment at all to the patient, who would then die. It does not follow that the surgeon’s performing the open-heart surgery does not harm the patient. Where some minor step – perhaps giving an aspirin – would have saved the patient’s life but the surgeon insists on performing open-heart surgery not to help the patient but to improve the surgeon’s own repertoire of skills, the patient *is* harmed by the open-heart surgery. Thus, to secure the “no harm done” result, we need to make the stronger assumption that the particular ART constituted the new person’s *only way* of attaining existence. Only on that stronger assumption does the claim that the ART does not harm the new child so long as the life conferred is worth living become plausible.

Similarly, the view Coleman attributes to others becomes plausible as well only on the understanding that the stronger assumption is in place. Thus:

As Bonnie Steinbock points out, the number of situations in which children can be considered harmed by being brought into existence “appears vanishingly small. Only for conditions which combine excruciating and unrelievable physical pain and such a brief life span that the child is unable to develop any compensating abilities could we assert the claim with any confidence.”¹⁹

And:

The difficulty of showing that children are harmed by being born with even severe impairments has led some commentators to conclude that the interests of future children will rarely provide a sufficient reason to object to the use of ARTs. According to John Robertson, for example, even if the children resulting from ARTs are likely to suffer tremendous disadvantages, “risking damage to offspring would not seem to wrong the offspring if it were not possible for them to be conceived or born without undergoing the risk of damage.” The only exception Robertson would recognize is when the risks to the child are so significant that the child's birth would satisfy the wrongful life standard – in other words, if “any life at all with the conditions of [the child's] birth would be so harmful to him that from his perspective he would prefer not to live.”²⁰

In the very passage Coleman cites, John Robertson explicitly puts the stronger assumption into place (thus Robertson writes that it is “*not possible for them to be conceived or born without undergoing the risk of damage*”).²¹ Plausibly, then, Coleman as well intends to limit his “no harm done” result to the case where that assumption holds.

In any event, it seems that it is a mistake to treat all ARTs alike in respect of harm. Some ARTs, and some procreative acts, harm the very offspring they serve to create and others do not. The very fact that an act is procreative does not immunize the agent who performs the act from the charge that what he or she or they have done *harms* the person that act brings into a worthwhile existence.

Nor does it appear that cases where other, safer routes into existence are available are all that rare. Supernumerary pregnancy is a major source of disability in the ARTs context.²² I have argued elsewhere that we can likewise identify safer routes into existence in the face of human reproductive cloning.²³

My view, then, is that each ART has to be evaluated on its own merits for harm done to each potential claimant. And the question in each case will be whether the agents have some way

of bringing that person into an existence that is *better for that person*. Asking that question in the context of ITISP, then, and focusing on any one of the group of six to eight babies who have made their way into existence via a supernumerary pregnancy with the scars to show for it, the answer as noted before is yes: for each such person, agents had the option of making things substantially better for that person by declining to bring all or most of the rest of that person's cohort into existence to begin with.

V. The Problem of Collective Harm

Consider a difficult scenario (Case I). The provider is all for avoiding ITISP. The patient, however, will consent and physically submit *only* to those treatments that *exclude* the relevant safeguards. She is perfectly willing to have the provider inject aggressive doses of ovulation-stimulating hormones. But she won't agree to egg retrieval and IVF coupled with a limitation on the number of eggs fertilized or embryos transferred to the uterus, nor to fetal reduction, nor to sonography for the purpose of determining the number of follicles stimulated, nor to any postponement of fertilization. The patient thus proclaims, accurately and truthfully, that "it's my way or the highway" for these particular, soon-to-be-damaged babies: she will produce them all or she will produce none.

The implications MCH1 generates in Case I are plausible but not complete. Thus, MCH1 implies that what the patient has done *as an individual* harms each of the damaged babies. It also implies that the provider *as an individual*, who, after all, lacks the option of creating additional wellbeing for any of the babies, all the better options having been blocked by the patient, has not harmed any of the babies. MCH1, finally, implies that what the agents have done *together as a*

group harms each of the babies. What MCH1 does not tell us, however, is what the harm done by the group means for each agent. Is the group-harm to be attributed to any or all members of the group, or not? That question may not seem to be pressing in the case of the patient who is independently recognized to have harmed each baby. It is, however, an important open question in the case of the provider.

What the provider has done clearly constitutes a significant part of the causal sequence of acts and events that eventuate in each baby's existing and suffering. The provider was, indeed, a knowledgeable, active (yet hardly essential or necessary, since other providers very likely would have allowed the patient to have had the babies "her way" had the one provider declined to do so) participant in bringing about that group-harm. On that basis, one might take the view that group-harm should be attributed to the provider, now as a participant in the group rather than on an individual basis.²⁴

In fact, however, the attribution of group-harm to the provider in the context of Case I seems inappropriate and indeed harsh.²⁵ While it perhaps gets at some truth or another – some truth, I think, about causation; the provider does seem properly characterized as having participated in bringing about the particular outcome for the baby – it trips over the "bigger" truth that this particular agent has *done the best that he or she could, given the options he or she had*, for the baby. The right thing to say about Case I thus seems to be that, while the provider and patient acting together could have done more for the baby than they have and have thus harmed that baby, the provider has not harmed the baby, either as an individual or as a participant in the group.

Consistent with this approach, we might believe that the concept of group-harm does not, after all, have any particular function to play when an agent imposes harm on a person. If, in the

end, we think that one person harming another is just a matter of whether that agent *as an individual* could have created more wellbeing for a given person than he or she has – a view that Case I examined in isolation may seem to suggest – then we need to add a condition to MCH1 that specifies that, if an agent as an individual has not harmed *p*, then the agent has not harmed *p* in any respect at all, including as participant in a group. We can then retain clause (2) of MCH1 – we have not, after all, seen any reason to think that it is *not* the case that the patient and provider have together as a group harmed the babies – but accept, at the same time, that a finding of group-harm under that clause (2) has no implications for whether a given agent has imposed harm in a given case.²⁶

In fact, downplaying the concept of group-harm in this way would be a mistake. Cases involving *reciprocally limited options* seem definitively to establish that the concept of group-harm has a central role to play in any plausible criterion of harm. Thus, in Case I, it is the patient who eliminates, from the perspective of the provider, certain better-for-the-baby options. But just as the patient may limit the provider's options, so may the provider limit the patient's. Consider, then, Case II, in which the two agents each claim, truthfully, that even if he or she had not opted to create less wellbeing for the baby rather than more the other surely would have done exactly that. On those facts, clause (1) of MCH1 implies that neither patient nor provider, as individual, has harmed the baby. Clause (2) implies group-harm, in virtue of the fact that the agents together as a group had the option of doing more for the baby than they have. Is it really plausible, given these facts, that neither patient nor provider has harmed the baby in any respect at all, including as participant? Is it really plausible, given these facts, that the *group* – that abstract object – has done

the dirty work, yet that dirty work cannot be attributed to *any* of the agents who constitute that group?

Cases more straightforward than ITISP similarly argue in favor of according a critical function to the concept of group-harm. Consider, for example, a variation (Case III) on Joel Feinberg's "two bullies in the schoolyard" example. In Feinberg's case, one bully, b1, beats the schoolboy; had he failed to do so, a second bully, b2, would have stepped in to do the job instead. On my variation, we stipulate not just that b2 *would* step in and beat the boy if b1 stepped back but also that no better-for-the-boy option exists for b1 than to beat the boy. Thus b1, for some reason or another, *can't*, for the purpose of giving the boy a chance (say) to flee, turn his thuggish strength against b2 rather than the boy. Let's stipulate, as well, that b1 and b2 are symmetrical in their relationships with the boy: just as b1 has no way of stopping b2 from beating the boy, neither does b2 have any way of stopping b1. To make the idea more palpable that the bullies' options are indeed so severely limited, let's further imagine that we are dealing with a shooting rather than a beating, with each bully effectively shielded from any change of heart on the part of the other and any chance that the other bully will turn his violence toward the one bully rather than the boy. Thus: b1 shoots the boy in the arm; and b2 would have shot the boy in the arm had b1 failed to do so.

On these facts, each of b1 and b2 has done his best for the boy, having selected, from all of the options that exist for the boy at the critical time, an option that is at least as good for the boy as any other. As individual, then, neither bully has harmed the boy, according to clause (1) of MCH1. If we now adopt the condition according to which "no harm done" by any agent *acting as individual* entails "no harm done" by anyone at all, we commit ourselves to the result that b1 has

not harmed the boy when b1 shoots the boy in the arm and that b2 has not harmed the boy when b2 eliminates b1's better-for-the-boy options. But those results seem clearly false. The boy has been harmed when he is shot, and it is b1 and b2 who have done the harming. That is so, despite the fact that b1 and b2 *as individuals* have harmed no one. Instead, by virtue of their participation in the group consisting of both b1 and b2, they have each, as participants in that group, harmed the boy.

The upshot is that, in the context of Case III and, I think, Case II as well it seems implausible to recognize that a group-harm has been imposed but to hold at the same time that, for each agent who is a member of that group, that agent has not harmed his or her apparent victim (whom the agent has, for example, just shot) in any respect at all, either as individual or as participant in the group. In fact, if ever the attribution of group-harm to members of the group is in order, it is in the context of Case III.

Still, just as Cases II and III suggest that the view that group-harm should *never* be attributed to the members of the group is a mistake, Case I suggests that the view that group-harm can *always* be attributed to the members of the group is a mistake as well. Our task, then, becomes that of deciding when group-harm is to be attributed to the agent and when it is not.

In completing this task, it will be useful to introduce one last case (Case IV). This case, yet another variation on the two bullies case, is analogous to Case I in important ways but has the advantage of simplicity. Suppose that b1 shoots the boy in the arm. Suppose, however, that the truth is that b1 is ready and willing to refrain from shooting the boy at all. What provokes b1 to shoot the boy in the arm is that b1 understands, correctly, that if he doesn't, b2 is fully committed to, and in fact will, step in and shoot the boy in the heart. b1's shooting the boy in the arm thus, in

effect, saves the boy's life. Attribution of the group-harm to b1 on these particular facts seems a mistake.

Under the plausible supposition that group-harm is to be attributed to b1 in Case III but not in Case IV, our task becomes to identify the critical difference between Case III and Case IV. But particularly striking are the many similarities between Case III and Case IV. Thus, in both cases, b1 and b2 acting together could have avoided shooting the boy at all; both cases, thus, represent instances of group-harm. And in both cases, b1 has done the best he as an individual can for the boy and thus plausibly has not, as an individual, harmed the boy. And in both cases, what b1 has done constitutes part of the causal sequence that eventuates in a certain outcome for the boy. But the basis for our attribution of group-harm to the agent cannot simply be the agent's participation in the causal sequence that eventuates in that group-harm. While such a test gives us the result we want in Cases II and III, it would also lead us to attribute group-harm to the hapless provider in Case I and the possibly heroic b1 in Case IV.²⁷ What, then, is the critical distinction that makes it plausible to attribute group-harm to b1 in Case III but not in Case IV?

At only one point do these two cases part ways. We have already noted that in Case IV b1 in effect *saves the boy's life*. In Case III, in contrast, the argument that b1 has *saved* the boy from anything (other than being shot in the arm by b2, which would have been no loss) falls flat. If the attribution of group-harm to b1 seems appropriate in Case III but not in Case IV, then it is this distinction between the two cases that we will need to exploit.

In spelling out exactly how it is that b1 "saves the boy's life" in Case IV but "saves" the boy from nothing of any moral significance in Case III, we are able to draw a distinction between the two cases. b1's participation in the group-harm in Case IV alone makes things better for the

boy rather than worse. Put another way, b1's *non-participation* in Case IV would have made things still worse for the boy. In contrast, b1's participation in the group-harm in Case III does nothing for the boy at all; and b1's *non-participation* – his opting out, in other words – would not have made things any *worse* for the boy.

The following revision of MCH1 incorporates these points:

Maximizing Criterion of Harm (second version) (MCH2)

Where p is any existing person, harm is imposed on p :

1. by an agent i acting as an *individual*, if i had the option of creating more wellbeing for p and instead created less;
2. by an agent i acting as a *participant* in the harming of p by a group of agents $i_1 \dots i_n$, if:
 - a. harm is imposed by agents $i_1 \dots i_n$ acting as a *group* (i.e., $i_1 \dots i_n$ together had the option of creating more wellbeing for p and instead created less);
 - b. i is a member of that group and i participates in the bringing about of that harm; and
 - c. i 's non-participation in the bringing about of that harm creates as much or more wellbeing for p as i 's participation in the bringing about of that harm creates for p (i.e., i 's participation does *not* create more wellbeing for p than i 's non-participation does).
3. An agent i who does not harm p either under clause (1) as individual or under clause (2) as a participant does not harm p at all.

Same notes as for MCH1, and in addition: "Together having the option of creating more wellbeing" does not imply coordinated acts or agents acting in collaboration but rather acts that in combination could have created more wellbeing. Similarly, " $i \dots$ participates \dots in the bringing about of that harm" does not imply coordination or collaboration but rather that i 's act (including omissions) is part of the causal sequence that in fact eventuates in a particular outcome for p . " i 's non-participation in the bringing about of that harm" occurs when i 's acts (including omissions) are not part of the causal sequence eventuating in a particular outcome for p .

It is clause (2)(c) of MCH2 that makes the desirable distinction between Case III and Case IV. In both cases, b1 had the option of *not* participating in what is done to the boy: b1 could have walked away in either case, leaving b2 to shoot the boy. However, in Case IV, that non-participation on the part of b1 would have made things worse for the boy by causing the boy to be shot in the heart; it is exactly on those grounds that we say that b1's participation makes things better for the boy and, indeed, saves the boy's life. Clause (2)(c) is thus not satisfied for b1 in Case IV, with the result that b1, acting as a participant, does not harm the boy. In contrast, in Case III, b1's non-participation – b1's walking away – would not have made things worse for the boy; correspondingly, b1's participation makes them no better, and we have no basis then for saying that b1 has rescued the boy from anything, even a beating. Clause (2)(c) is thus satisfied for b1 in Case III.

An intuitive way of putting this point is just to say that b1 has no good reason, or excuse, in terms of the boy's own wellbeing, to remain in the fray in Case III but every good reason to do so in Case IV.²⁸ On the assumption that harming someone or another is a necessary condition for wrongdoing, then, we can say that MCH2 gets b1 off the moral hook in Case IV but not Case III.

Given the analogies between Case I and Case IV, we also go a long way toward getting the health care provider off the moral hook in Case I as well. Thus, where the provider is ready and willing to put safeguards against supernumerary pregnancy into place but the patient is hopelessly and demonstrably recalcitrant, and where the ITISP-damaged babies in fact exist and suffer, MCH2 provides a defense against the claim that the provider, acting either as an individual or as a participant, has harmed those babies.

There is, however, more to be said about Case I. In real life, the provider may well find it difficult to establish that his or her non-participation in the causal sequence that resulted in the supernumerary pregnancy in fact would have led to the ITISP-damaged babies never having existed at all and thus have made things worse for each of them.

The provider's defense under clause (2)(c) of MCH2 in the context of Case I – and, in legal terms, we are here talking about an affirmative defense, the kind the provider ideally pleads and certainly must prove – is at best a murky one. Its success requires a showing that the provider *lacked the option of swaying* the patient to accede to some more reasonable treatment plan, one that included the relevant safeguards. Only on such a showing can we conclude that the provider's non-participation would have made things worse for each of the damaged babies. Is it really so clear that, had the provider specified, as a condition of treatment, various safeguards against supernumerary pregnancy, the patient, however reluctantly, would not have accepted that condition?²⁹

For purposes here, we make the assumption that the patient has consented to the risks to her own health created by ITISP. (She cannot, of course, effectively consent to the risks that ITISP creates for her future babies' health.) But that consent is not relevant to whether the patient, if pressed, would have also consented to a treatment plan more protective of her future offspring. Thus, the patient may with perfect consistency consent to the risks to her health yet after-the-fact quite truthfully report that she was not so firmly or intractably committed to the “her babies, her way” notion that she would not have accepted the provider's treatment plan had the provider insisted on limiting the patient's options rather than the other way around.

In Case IV, b1 faces a similarly murky evidentiary basis. When wading into the business of saving lives by way of shooting people in their arms, b1's defense may lose its credibility in the face of b2's own story that, had b1 put down his gun and refrained from shooting the boy in the arm, so b2 would likely have put his gun down as well, notwithstanding his "tough talk," and refrained from shooting the boy in the heart. The guesswork involved in grasping what "would have been" is never going to create an entirely comfortable evidentiary base for the kind of defense clause (2)(c) provides.

This is not to say we cannot imagine cases in which it is entirely plausible and credible that the patient *cannot* be swayed to a more reasonable treatment plan. The clearest example of such a case is where the patient cannot *afford* any such plan. Perhaps she would be all for safeguards against ITISP but is barely able to pay for even one ovulation-stimulating treatment – and so she wants it to be an aggressive one. From the patient's point of view, then, safeguards against ITISP – and in particular identity-preserving safeguards, which may be more expensive – simply do not constitute an option. But that the option does not exist for the *patient* does not imply that it does not exist for the *provider*. In many cases, where the patient cannot pay, the option exists for the provider to forgive or foot the bill or otherwise arrange for the service to be covered. According to MCH2, the very existence of such a far-out, "provider pays," identity-preserving option implies that the ITISP-damaged baby has been harmed by the provider – acting as an individual now, under clause (1) of MCH2, and not simply as a participant in a group.³⁰

This result is particularly plausible given that the existence of the "provider pays" option does nothing more under MCH2 than establish *harm*. It thus does not even begin to suggest that the provider has an *obligation* to forgive or foot the bill for expensive, identity-preserving

safeguards when a patient is willing to make use of but unable to pay for those safeguards. Nor do we obtain any such suggestion from the broader, person-based moral theory that MCH2 fits into. We would, instead, obtain the more modest result that, on the assumption the provider has the option of doing so, he or she must insure that any babies who are brought into existence are left undamaged by ITISP. Consistent with that more modest result, the provider – the doctor; the clinic; the insurance company – remains perfectly free to adopt cheaper safeguards against ITISP – sonography, e.g., followed by instructions to delay fertilization until the next cycle alongside screening to insure the patient’s willingness and ability to follow such instructions – instead of any more expensive, identity-preserving option. What the provider cannot do is help to produce the babies but at the same time fail to put *any* safeguards – identity-preserving or not – against ITISP in place in the case where the provider has the option of doing so.

The standard imposed on the provider by this account may seem too stringent. But it should be kept in mind that the provider always yet another option – short of putting safeguards against ITISP, whether identity-preserving or not, into place – for avoiding harm to the babies he or she helps to bring into existence. Non-participation, in particular, always remains an option for the provider. The provider may always eliminate himself or herself from the causal sequence that ends in the coming into existence of the ITISP-damaged babies. In that case, the provider, though remaining a member of the group consisting of both provider and patient, does not *participate* in the imposition of a group-harm and thus does not satisfy clause (2)(b) of MCH2. Suppose, then, that the patient proclaims that she will under no circumstances consent to safeguards against ITISP. But suppose too that the provider is unsure how intractable that sentiment really is, or at least is unsure that he or she will be able to demonstrate, in the face of a moral or legal charge, that

the patient's mind could not have been changed. The provider always has the option of not accepting that particular patient for that particular service and instead watching, as it were, from afar, removed from the fray, as the patient works to convince some other fertility specialist to administer the aggressive fertility treatments that, in the absence of safeguards, bring the ITISP-damaged babies into their troubled existence. The provider who eliminates himself from the causal sequence in this way – who, for example, requires that safeguards against supernumerary pregnancy be put in place as a condition of treatment, knowing, or guessing, in advance that the patient will not agree to that condition – will not succeed in improving the lot of any of the babies. But the provider will have insured that he or she is not the one who has done the patient's dirty work. Whatever the details of our full moral account, that fact is likely to be deemed a highly significant one.

Thus, the defense that clause 2(c) of MCH2 makes available to the health care provider – “I did not harm but rather rescued those babies from the plight of having never existed at all” – is lovely in theory. But it may often lack significant practical value to providers against whom claims are brought in the ITISP context.

VI. The “Probabilistic” Version of the Nonidentity Problem; the Nonidentity Fallacy

We have already examined, and discarded, one type of nonidentity problem as it arises in the context of ITISP. Thus the argument was made, in Part III, that ITISP is nothing like wrongful life. But a still more deadly form of the nonidentity problem is one that suggests that in assessing harm the critical question is whether, among the various facially better-for-the-baby options that exist for the agent at the critical time, any such options create any *real chance* for that baby to

come into existence at all. If the baby's chances of existence under each such seemingly better-for-the-baby option are practically zero, then it is hard to see how any of those options could *really* be better-for-the-baby than what has in fact been done for the baby. How can the virtual certainty of never having existed at all be better for the baby than a life that is, though flawed, unambiguously worth living?³¹

Now, the idea that the options – which I have labeled “identity-preserving” – I have described for how the agents may have made things better for the very same, identical ITISP-damaged babies in fact would have taken each and every such baby off-track for existence altogether may seem, on first inspection, a compelling one. It is an idea that is reflected in Gregory Kavka's discussion of the “precariousness of existence.”³² When it comes to the matter of existence, timing is everything: a conception that takes place a month earlier or later, or a minute earlier or later, will very, very likely yield a different child altogether. Make even the most minute change in the sequence of acts and events that eventuates in a given person's coming into existence and we very, very likely end up bringing someone into existence who is distinct from – “nonidentical” to – the original.³³

How does the phenomenon of the “precariousness of existence” apply in the context of ITISP? Suppose that a particular supernumerary pregnancy occurred when fertility drugs were used in an aggressive way, and suppose insemination then takes place by way of sexual intercourse. Eight babies, say, are conceived and then damaged over the course of the next several weeks and are finally born with an array of serious disabilities. I argued in Part III above that, in virtue of the fact that agents had the option for doing better for each one – though not all – of those eight babies than they have, each one has been harmed by what the agents have in fact opted to do.

But suppose that the agents, instead of acting as they did, put one of my prescribed “identity-preserving” safeguards into place. Suppose, indeed, that they *tried with all their might* to bring one or two of the eight damaged babies into a better existence by way of leaving the remainder of the cohort out of existence altogether. Instead, then, of insemination by sexual intercourse, the eight – ten? twelve? – ripened eggs would have been retrieved via laparoscopy and inseminated using the father’s sperm. Then, two or three of the embryos would have been transferred to the uterus, with the remaining embryos being either frozen or discarded. Aren’t the chances that the *same* sperm cells collected at a different time and by a different process – out of all those hundreds of millions of sperm cells that may have been produced over a period of a few days – would end up inseminating the *same* egg cells vanishingly small? The one or two babies then born would have been healthier babies, to be sure, but surely the odds are overwhelming that they would have been distinct babies from – “nonidentical” to – any of the original, damaged babies. And for each damaged baby in fact born, surely it is better to have the flawed existence than it is to be allocated only a one-in-a-zillion chance of ever having existed at all.

This “no harm done” argument is beguiling – and, I have argued elsewhere, fallacious.³⁴ The argument is motivated by the observation that the flawed existence is surely better for the child than is the very small chance of ever having existed at all. But betterness determinations – determinations, in other words, that one act is better for p than another, in the sense that it will produce a better outcome for p than the other – are inherently tricky, and one apparently correct principle is that betterness determinations are not reliable when based on a comparison between *actual* wellbeing – the wellbeing level p in fact ends up with, given the performance of the one act – and *expected* wellbeing – a calculation based on the summation of the wellbeing levels p has in

each possible outcome multiplied by the probability, for each outcome, that that outcome will occur, given the performance of the other act.

The fallacy is deeper than this, however. After all, it is easy to shift the terms of the argument to avoid the problematic actual-against-expected comparison and proceed instead from a more legitimate expected-against-expected comparison. The idea then becomes that, surely, p 's chances for existence are at least improved when insemination takes place naturally, as it in fact does, rather than artificially, in the identity-preserving manner I have described.

In the face of this slightly revised form of the nonidentity problem, however, it is important to keep in mind the fact that probabilities – at least conceived in a way that will make them useful in guiding the choices we make; that is, as determined by what agents grasp or are in a position to grasp at the critical time *just prior to choice* – are themselves in flux, changing in dramatic ways as the future unfolds. Thus, let t be that moment just prior to action (at any time after that, the choice will already have been made; at any earlier time, agents will not yet be in their epistemically best position). At t , it is within the agents' control to choose natural insemination. But at t there exists a vast number of possible ways and times at which natural insemination may itself take place. Admittedly, if the agents had opted for *artificial* insemination instead, millions of possible sequences of acts and events, leading to millions of possible babies, would have stretched out before the agents from time t , and the chances of any one baby's coming into existence as of t would have been quite small. But so what? Precisely the same point holds – at t – for *natural* insemination. The chances of any one baby's coming into existence, at t , is quite small in the case of natural insemination as well.

The upshot is that, after all, doing things as the agents have offers, at the critical time, no advantage to any of the ITISP-damaged babies over doing things in a way that would have safeguarded the wellbeing of that same baby. The phenomenon of the precariousness of existence thus cuts both ways: existence, for any given damaged baby, is highly precarious, independent of whether the agents opt for natural or artificial insemination. But the outcome, if the baby does make it into existence, would have been much better for the baby had the agents put safeguards against supernumerary pregnancy into place. Expected wellbeing for each damaged baby is thus, after all, greater if agents take the safer course than if they take the riskier course. The upshot is that the argument that the riskier course is somehow better for each of the babies in fact born, and thus cannot harm them, evaporates.³⁵

VII. Conclusion

My argument is that, while harm may seem elusive in the ITISP context, a plausible criterion of harm and a close inspection of the facts of the case, including all relevant probabilities, allows us to discern harm and to challenge arguments that try to establish, in the face of babies born blind, deaf, palsied and disabled for life, “no harm done.” In the end, our intuitive and, I think, person-based, judgments about these grim cases seem to align with the better argument and the better analysis. Now, we cannot always trust our intuitive assessments. But they may sometimes function quite nicely as a red flag. And that is just how they function in the case at hand, where they alert us to the defects in a potpourri of otherwise beguiling arguments to the false “no harm done” result.

*M.A. Roberts completed the J.D. at the University of Texas at Austin and a Ph.D. in philosophy at the University of Massachusetts at Amherst (Five College Ph.D. Program). Roberts writes in the area of law and ethics, focusing especially on harm and the structure of morality.

An earlier version of this paper was presented as part of the Hart Lecture Series at the University of Pennsylvania Center for Bioethics (June 2006). The author is extremely grateful for the comments of Arthur Caplan, Vardit Ravitsky, Paul Wolpe and other participants in that seminar, as well as for comments received on prior drafts of this paper from Alan McMichael and David Wasserman.

References

1. “The greatest harm to the children of ART is posed by the probability that they will be born a multiple. . . . Approximately thirty-six percent of all births resulting from [in vitro fertilization] and [intracytoplasmic sperm injection] are multiples . . . and the numbers continue to grow. . . . Multiples generally are considered a necessary side effect of existing ART practices It is well-established that multiple births cause serious health problems. Multiples are more likely to be born prematurely . . . and at [low birth weight] The mortality rates for multiples are significantly higher than for singletons: for triplets, it is thirteen times that of singletons, and for twins, it is five times . . . higher. . . . Even if they survive, multiples (including twins) are more likely to suffer serious physical and mental handicaps, such as cerebral palsy. . . . In light of this evidence, one noted commentator has wondered why the widespread creation of multiples through ART has not been considered a ‘public health problem.’” J. L. Rosato, “The Children of ART (Assisted Reproductive Technology): Should the Law Protect Them From Harm?” *Utah Law Review* 57 (2004): 77-78. S. Khanijou notes that “[i]n the last decade, there has been a dramatic increase in multifetal pregnancies as a result of ART. . . . Data published by the Centers for Disease Control . . . indicate that more than 30% of assisted technology births are multiples compared to the 2% incidence in the general population. . . . The media and much of the public hail multiple births as a ‘miracle,’ as in the case of the septuplets born to the McCaughey family, but the medical community is not so enamored. They have identified such pregnancies as failures, rather than successes, of the IVF enterprise. The press heralded the live birth of sextuplets in California, but the attention soon disappeared when three of the six died. In 1998, one year after the septuplets, the birth of octuplets in Texas made history. But this holiday ‘miracle’ soon turned tragic when one of the infants died and the remaining seven returned to neonatal intensive care soon after Christmas. Perhaps some records were simply not meant to be broken.” See S. Khanijou, “Multifetal Pregnancy Reduction in Assisted Reproductive Technologies: A License to Kill?” *DePaul Journal of Health Care Law* 8 (2005): at 403, 404. See also N. S. Green, “Risks of Birth Defects and Other Adverse Outcomes Associated With Assisted Reproductive Technology,” *Pediatrics* 114 (2004): 256-59.

2. The “no harm done” result turns on a showing that an impaired existence is the *only* way that a particular child may exist at all. As one court put the point, wrongful life constitutes “a Hobson’s choice [between] life in an impaired state and nonexistence.” *Becker v. Schwartz*, 386

N.E.2d 807, 812 (N.Y. 1978). Notably, the standard “no harm done” wrongful life model may not even be applicable to all instances of wrongful life. Where the child who is negligently brought into existence foreseeably experiences so much pain and misery and so little of whatever it is that we think makes life worth living, it is plausible to say that the child would have been better off never having been born at all and thus is harmed when he or she is brought into existence. See note 11 below.

3. For discussion of the risks of congenital and chromosomal abnormalities and developmental delays associated with ICSI, see Rosato, *supra* note 1, at 69, notes 86-88. The criterion of harm set forth in Part III below and elaborated in Part V below implies that, in the case where the agents had *no option* for bringing the impaired child into existence other than by way of ICSI, that child is not harmed by ICSI despite the fact that ICSI may be a part of the causal sequence that results in impairment to the child. Thus, according to that criterion, where the impairment that is imposed is a *necessary condition* for the child’s very existence, and where that existence is worth having (i.e., at or above the zero level implied by nonexistence, and hence not genuinely “wrongful”), the use of ICSI does not harm the child. Of course, consistent with this result, if agents had available a safer ICSI technique, or the option of substituting the ICSI technique in fact used with some safer infertility treatment, such that the agents could have brought the damaged child into existence without the damage, we can take the position that the form of ICSI in fact used *harms* the child. The position recently articulated by John A. Robertson seems roughly consistent with this point. See J. Robertson, “Procreative Liberty and Harm to Offspring in Assisted Reproduction,” *American Journal of Law and Medicine* 30, no. 7 (2004): 14-15. (“In all of the situations surveyed, the child appears to be harmed by the very method of conception, gestation, or social setting of birth. If, as many commentators argue and legislative schemes provide, ‘the prime consideration must be the welfare of the child,’ then use of ARTs in those situations described must be questioned. To prevent the feared ‘injury’ to the child, the person and couple should give up the use of the ARTs that pose those risks. . . . But this leads to a paradox. If offspring are ‘harmed’ by being born in those conditions, then the only way to prevent the harm is to not use those techniques. But this means that the children sought to be protected will never be born. Because their lives will not be so miserable as to be ‘wrongful,’ it would seem that once born they have benefited from rather than been harmed by being born. If that is so, then using ARTs to enable their birth does not harm them and does not justify restriction on those grounds. . . . The logic of this position troubles many people. . . . A key point about the paradox of non-wrongful life is that the person could not have been born without the condition of concern. If so, refusing the act or omission that causes the child to be born with that condition cannot harm the child. Of course, if changes in technique or treatment protocols could reduce the frequency of the condition, there would be an obligation to adopt those changes. However, in situations in which no improvement can be made one cannot show that the child has been harmed as a result.”)

4. I track Jan Narveson’s language here, who writes that “[w]e are in favour of making people happy,” and “neutral about making happy people.” See J. Narveson, “Moral Problems of Population,” in *Ethics and Population*, Michael D. Bayles, ed. (Cambridge, MA: Schenkman, 1976): at 73.

5. D. Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1987): at 363 (emphasis eliminated in part).

6. Government regulation in this area can be through the existing tort law framework or through direct regulation or both. “[T]wo visions of federal regulation have been proposed: one by The President's Council on Bioethics (‘Council’); and another by Erik Parens and Lori Knowles of The Hastings Center. Although both proposals are extremely thoughtful, they go further than necessary and ultimately will lead to overregulation of ART.” Rosato, *supra* note 1, at 90. For the argument that the FDA should restrict or even withdraw certain drugs used to stimulate ovulation from the market, see L. Noah, “Assisted Reproductive Technologies and the Pitfalls of Unregulated Biomedical Innovation,” *Florida Law Review* 55 (2003): at 603. For further discussion of the constitutional limitations on the government’s ability to regulate conduct in the context of supernumerary pregnancy, see Rosato, *supra* note 1.

7. We might, for example, eschew the highly person-based “harm principle” associated with John Stuart Mill’s *On Liberty* in favor of an affiliational or a communitarian model of the U.S. Constitution and the state constitutions, and, indeed, of the role of law in society. More narrowly, we might revise constitutional privacy law and the law of negligence to include “Principle N” developed by Buchanan, et al., or Parfit’s own same-numbers principle “Q.” According to Principle N, agents are morally obligated to take steps to insure that any new person they produce will have an existence that avoids “serious suffering or limited opportunity or serious loss of happiness or good” in any case in which the agents could have produced a better-off, though “nonidentical,” child in place of the impaired child without imposing “substantial burdens” on anyone, including the agents themselves. A. Buchanan, D. W. Brock, N. Daniels and D. Wikler, *From Chance to Choice: Genetics and Justice* (Cambridge University Press, 2000): at 249. Parfit’s Q suggests that it is morally better to bring a better-off person into existence rather than a worse-off person. See Parfit, *Reasons and Persons*, *supra* note 5, at 369. For a critical discussion of these approaches, see, among other sources, Robertson, *supra* note 3, at 15-16.

8. This is so, whether the impersonal approach under scrutiny is a traditional form of utilitarianism (sometimes called “totalism”) or a form of consequentialism that considers the overall good to be constituted not by aggregate wellbeing alone but by other values as well, such as equality (“pluralism”). The repugnant conclusion, the infinite population problem and an array of extreme inequality problems clearly challenge totalism. While pluralism avoids some variations on each of those problems, other variations are more difficult for the pluralist to provide a plausible account of. The very case at hand illustrates potential problems with a move to totalism. Totalism, applied to ITISP, seems to generate, depending on the specific facts of the case, what is in effect a *mini-repugnant conclusion*: the result that it is perfectly fine, indeed, obligatory, to continue, or even to commence, a supernumerary pregnancy in order to bring six, seven, or eight babies into existence in one fell swoop rather than just a single healthy baby, on the grounds that the former option is the one that increases total wellbeing. For additional discussion of these points, see M.A. Roberts, “Supernumerary Pregnancy and the Limits of the Constitutional Privacy

Guaranty," *Jo. of Philosophical Research*, ed. F. Adams (special supplement) (Philosophy Documentation Center, 2005) 105-117, and "A New Way of Doing the Best That We Can: Person-Based Consequentialism and the Equality Problem," *Ethics* 112, no. 2 (2002) 315-50.

9. Such criteria fall into one of two classes: "adversity-based" approaches to harm and "substitutional" approaches to harm. Adversity-based criteria discern harm in any cases involving substantial pain or hardship, even when that pain or hardship is itself physically necessary to secure some clear good for the very individual who suffers the pain or hardship. Such criteria thus by design overlook the fact that acts that impose pain or hardship may in some cases also happen to be the very acts that maximize wellbeing. The adversity-based approach basically equates harm with something like pain: an unpleasant or miserable experience that is not in itself a plus but that we immediately understand to be something we may very badly want simply because, all things considered, it's good for us. Here we might point to Rosato's apparent identification of harm with any serious physical or emotional adversity, or Elizabeth Harman's explicit identification of harm with "significant bodily damage" – including, e.g., the damage that is done when the surgeon makes an incision in one's abdomen for the purpose of saving one's life through the removal of a ruptured appendix. Thus Harman writes: "I claim that *causing* pain, early death, bodily damage or deformation is harming." E. Harman, "Can We Harm and Benefit in Creating?" *Philosophical Perspectives (Ethics)* 18 (2004): at 92, 93. "[T]he fact that adopting the Policy would harm her [in this sense] is a reason against adopting the Policy," at 92. "[A]dopting the policy is . . . wrong because we have an alternative that would involve parallel benefits, without parallel harms," at 93. Substitutional criteria provide that (with certain qualifications) harm has been done to one person just in case agents could have brought a distinct ("nonidentical") better off person into existence in place of that one. Philip G. Peters, Jr. has developed a substitutional approach to harm. P. G. Peters, Jr., *How Safe is Safe Enough?* (Oxford University Press, 2005). As Robertson observes, however, Peters' approach "only discerns harm in a case where the agents have the option of substituting a better off person in for a less well-off person. In the context of fertility treatment, however, that precondition may very well not be satisfied. . . . The [same number] principles that Brock and Parfit enunciate are attractive, and can be adopted without accepting a full-scale utilitarianism in all areas of life. But the need to keep the numbers the same and not unreasonably burden parents in making substitutions constrains its impact." Robertson, *supra* note 3, at 16.

10. For brief sketch of Peters' substitutional theory of harm, see note 9 above.

11. Consistent with the view that the provider's negligent act does not impose a *procreative* harm on the child who owes his or her very existence to that negligent act, we may think that that same act imposes *distributive* harms on persons *other* than the child and, perhaps, on the *child* as well. Thus, while the *procreative* harm may itself be unavoidable if the child is to exist at all, that same proposition does not hold for the various *distributive* harms that may be separately imposed on, for example, the child's siblings as well as the child himself or herself, depending on how the parents go about balancing their own scarce resources among family members. Distributive harm to the *parents* is recognized by almost all courts as creating the basis for an action for "wrongful birth." The suggestion here, however, is that such distributive harm to

both the child's existing and future siblings as well as to the child himself or herself may also create the basis for a cause of action. See, e.g., M.A. Roberts, "From Chance To Choice To Harm To Persons: Why We Should Favor a Person-Affecting Account of Wrongful Disability" (unpublished ms.).

A second point is important in this context. Courts may be correct in holding that *most* wrongful life cases – in which the child is born, say, deaf, or perhaps with Down syndrome – as a matter of empirical fact do not involve procreative harm to the child. However, it seems a mistake for any court to hold that, as a matter of law, the case in which the existence is so full of pain and misery as to itself constitute a harm *cannot* arise. Surely, after all, we understand that, at least in rare cases, never having existed at all will constitute the better alternative from the perspective of the individual who exists and suffers. For further discussion of the cogency of the wrongful life claim, see "Can It Ever Have Been Better Never to Have Existed At All? Person-Based Consequentialism and a New Repugnant Conclusion," *Jo. of Applied Philosophy* 20, no. 2 (2003) 159-85.

12. The federal district court in *Morgan v. Christman*, a case discussed aptly by Lars Noah in a recent paper, thus seemed to get this point just exactly right. See Noah, *supra* note 6, at 603.

13. At present, doctrines of intra-family immunity suggest that that patient will not be held liable for the harm she has done her offspring. At the same time, laws against parental abuse and neglect may be applicable. What account of harm we ultimately adopt will be of critical importance to whether doctrines of intra-family immunity require revision and just how laws against parental abuse and neglect apply in the ITISP context.

14. In part to capture the causal element in harming, we may require for harm both an increased risk as well as a flawed outcome. If so, then MCH1 could be revised to incorporate not just the concept of *actual* wellbeing but also the concept of *expected* wellbeing.

15. This is the approximate description of the plight of the Chukwu octuplets. For a description of that case, see R. Chester, "Double Trouble: Legal Solutions to the Medical Problems of Unconsented Sperm Harvesting and Drug-Induced Multiple Pregnancies," *St. Louis Law Journal* 44 (2000): at 451, 462-63. See also Roberts, "Supernumerary Pregnancy and the Limits of the Constitutional Privacy Guaranty," *supra* note 8.

16. I am grateful to David Wasserman and Arthur Caplan, respectively, for these points.

17. My suggestion here, then, is that the person-based consequentialist might appeal to the principle of "leximin" – the idea, that is, that in a trade-off situation, the agent's obligation will be first to increase the wellbeing of the least well-off, then to increase the wellbeing of the next least well-off, etc. For further discussion of this point, see M.A. Roberts, "A New Way of Doing the Best That We Can: Person-Based Consequentialism and the Equality Problem," *supra* note 8.

18. Carl H. Coleman, "Conceiving Harm: Disability Discrimination in Assisted Reproductive Technologies," *UCLA Law Review* 50 (2002): at 44-45.

19. Coleman, *supra* note 18, at 46-47 (citing Bonnie Steinbock, *Harm and Future Persons*, 2000 (unpublished manuscript)); see also B. Steinbock and R. McClamrock, "When Is Birth Unfair to the Child?" *Hastings Center Report* 24 (1994): at 15, 16.

20. Coleman, *supra* note 18, at 44-47.

21. Emphasis added. Consistent with this point, Robertson elsewhere acknowledges that, even outside the case where the life is less than worth living, harm to offspring may be imposed by ARTs in cases where agents have some way of "improving welfare" for offspring – some way, that is, of avoiding imposing "harm" on them while still bringing those same persons into existence. In this context, he cites the case of "multiple gestation." Robertson, *supra* note 3. Robertson then goes on to say that, "[u]nfortunately, few ART situations will be covered by the principle of avoidability."

22. *Supra* note 1.

23. See *Child Versus Childmaker: Future Persons and Present Duties in Ethics and the Law* (Rowman & Littlefield, 1998): at 179-211 (Chap. 5); "Cloning and Harming: Children, Future Persons and the 'Best Interest' Test," *Notre Dame Jo. of Law, Ethics & Public Policy* 13, no. 1 (1999): at 37-61; and "Human Cloning: A Case of No Harm Done?" *Jo. of Medicine and Philosophy* 21, no. 5 (1996): at 537-54. If the man-made proliferation of my genome counts as a diminution in wellbeing for me, that is, a *harm* to me – and I think that it is must, else we would not insist on the right of consent to the use of our genomes in connection with the production of others via cloning – then we can explain why cloning harms the very people it helps to bring into an existence worth having. Suppose, for example, that four of us have been produced by way of cloning – to keep things simple, let's say by way of embryonic cloning. By logic analogous to that used in Part III of this present paper, agents had the option of creating additional wellbeing for me – that is, the option of bringing me into a still better existence. All they had to do was bring me into existence and refrain from bringing the remainder of my cohort into existence. Rather than transferring all four embryos to one or more uteruses for development, they could have, in other words, transferred just the single embryo, the one who becomes me, and let the others perish. But that means they had the option of creating more wellbeing for me when they instead created less. And that means harm.

24. I treat here an issue of harm. But issues of causation arise as well. Michael D. Green discusses the intersection of these issues. Plausibly, what the provider has done at least should be counted as a "cause" of the outcome in which the baby both exists and suffers. At least, the provider's acts would seem to satisfy the NESS (the "necessary element of a sufficient set")

criterion for causation described by Green as having been developed by Wright and based on the work of Hart and Honore. Since the patient might always have gone to another provider, however, those same acts would seem not to satisfy the traditional “but for” test for causation. M. D. Green, “The Intersection of Factual Causation and Damages,” *DePaul Law Review* 55 (2006): at 671, 683 (citing R. W. Wright, “Causation in Tort Law,” *California Law Review* 73 (1985): 1735, 1795; and H. L. A. Hart and T. Honore, *Causation in the Law* (2nd ed.) (Oxford University Press, 1985): at 239).

25. “Cause the tortfeasor’s act is, but liability should not follow.” Green, *supra* note 24, at 700.

26. Parfit raised this question, though in a different context, in a recent talk entitled “What We Together Do: Questions About Collective Responsibility,” University Center for Human Values, DeCamp Lecture (Princeton University, October 10, 2005).

27. The Restatements of Torts describe classic cases involving reciprocally limited options as well as the broader problem of causal overdetermination (or multiple sufficient causes). In one such case, two fires are negligently set. Each fire is sufficient to destroy the owner’s property. The fires then converge on the owner’s property, destroying it. In that case, the agents have, by hypothesis, each engaged in a negligent act; they have, that is, acted (or failed to act) in a way that is deemed to breach a duty of care they owed to the property owner. A variation on the two fires case involves two agents negligently firing their guns and hitting the same eye of the same person at the same moment. Case III, in contrast, does not involve a double shooting. Rather b2 simply stands to the side, prepared to shoot if b1 doesn’t, as b1 in fact does the shooting. But it is not b2’s simply standing there that makes Case III an interesting one, nor his churlish failure to rescue (he can’t). It is rather the fact that – like the second fire in the two fires case – he is limiting the options that b1 would otherwise have had to just one: the boy’s being shot. Michael D. Green describes the view expressed in the Restatement (Third) of Torts: Liability for Physical Harm 27 cmt. i (Proposed Final Draft No. 1, 2005)(William C. Powers, Jr. and Michael D. Green, reporters): “If two hunters fire negligently at their quarry and their bullets arrive in the plaintiff’s eye at the same moment, is either one a cause of the plaintiff’s harm? This is just a variation on the familiar ‘two fires’ hypothetical that requires some consideration of the role of the sine qua non test for factual cause. Each hunter is a cause of the harm and liable for the full extent of damages suffered by the victim. *Sine qua non* will not do here, but there is no doubt that instinctively, rationally, and normatively, we are entirely comfortable describing each hunter’s action as a cause of the victim’s harm and holding each liable for damages affiliated with the loss of an eye.” Green, *supra* note 24, at 681. Green notes that the draft Restatement takes a similar position in the case where the overdetermination is due not to two tortious acts but rather to one tortious and one innocent act: “In conclusion, causation has taken the day and determined that the tortfeasor who concurs with an innocent cause is liable for damages reflecting the overdetermined harm. Note that the audience is no worse off because of Landsman’s tortious (and heinous) act, as they would have suffered equally in a tort-free world.” Green, at 687. Thus, the relevant section “applies regardless of whether the other cause involves tortious conduct or consists only of

innocent conduct. As long as the competing cause was sufficient to produce the same harm as the defendant's tortious conduct, this Section is applicable.” Green, at 687 (citing draft Restatement at 27 cmt. d). MCH2, below, similarly addresses the case of the one tortious and one innocent act in the same way it addresses the case of the two tortious acts. Another challenging scenario arises when both an agent’s act and some natural event are sufficient to impose a diminution in wellbeing for *p*. Do we still say that the agent has harmed *p*? MCH2, as drafted, will not have that implication. We are left then, under MCH2, with the result that the agent may be considered to harm *p* when another agent’s innocent act is sufficient to cause that same harm but not when nature is sufficient to cause that same harm.

28. I note that the interesting rule Joseph Mendola suggests for when agents are entitled to defect from a beneficent group project parallels MCH2 in certain respects. “Consequentialism, Group Acts, and Trolleys,” *Pacific Philosophical Quarterly* 86, no. 1 (2005): 64-87. Cases analogous to Case IV abound. For example, do slaves who lack any ability to escape their plight harm their own child when they bring that child into existence as a slave? MCH2, plausibly, suggests that they do not. Since the slave-parents, in producing the child as a slave, have thereby done as much for their child as they can, they have not, acting as individuals, harmed that child under clause (1) of MCH2. The issue of group-harm is somewhat more complex. It seems that a group-harm has been imposed on the child, given that the group consisting of the parents together with the various agents who maintain the institution of slavery had the ability to make things better for the slave-child than they are. But in this particular case the parents’ *non-participation* in the causal sequence that ends in the child’s existence as a slave would not have created as much wellbeing for that child as he or she himself in fact enjoys. The parents’ non-participation, in other words, would have made things still worse for the child. (We here make the assumption that the particular existence, despite its flaws, is better, from the child’s own point of view, than never having existed at all would have been.) Since what the parents have done thus does not satisfy clause (2)(c) of MCH2, we infer that the parents have not, acting now as participants in the group-harm, harmed that child. Clause (3) of MCH2 then provides that the parents, having harmed the child neither as individuals nor as participants in the group-harm, have not harmed the child at all. Environmental cases – cases in which parents bring children into an environment polluted by agents acting in the past and outside the influence of the parents – can be given an analogous account.

29. In real life, providers do have more sway over what the patient decides, and more options, than b1 has with respect to b2. Suppose that the patient informs the provider that she will not undergo any treatment that includes any safeguard against supernumerary pregnancy. And suppose her reasons are religious, or “ethical,” in nature, and have nothing to do with costs: she just thinks that once those egg cells pop out and begin to drift downwards in the direction of the fallopian tube, ripe and ready for fertilization, the “be fertile and populate the earth” principle or perhaps the still more important and morally weighty “more the merrier” principle is triggered. The provider informs the patient, then, that that’s not how he (suppose) or his clinic does business: no fertility treatment can be made available in the absence of precautions to avoid supernumerary pregnancy. And he explains to the patient, in the most graphic terms, why that is the case. Is this

option – clearly available to the provider – one that really so obviously would have failed to create additional wellbeing for the disabled child? Isn't the woman highly committed to having a child just as likely to accede to the provider's terms? Perhaps, by taking a firm stand, the provider does reduce in some small way the child's chances of existence – but surely many supernumerary pregnancies themselves are spontaneously aborted and surely a more limited pregnancy has a greater likelihood of success, so there may be something of a wash here. Especially in the case where after the fact, in the context of a civil action, the woman then asserts that, had the provider made precautions against supernumerary pregnancy a condition of treatment, she would have acceded to that condition, despite her vague religious or moral protestations to the contrary, I am not sure that we would not or should not find her testimony credible.

30. Notably, the option is not really so far-out. Imagine the reaction if a pediatrician charged \$100 for stitches administered under a sterile procedure but only \$50 for stitches administered otherwise. Do we really think that if the wound then becomes seriously infected the pediatrician has not *harmed* the child? Don't we rather think that the pediatrician is welcome to charge \$50 for the service rather than turn the child away but, having accepted the child as patient, must now clean the wound before stitching it?

31. I discuss this particular type of nonidentity problem – the “can't-expect-better” problem – elsewhere. See “The Nonidentity Fallacy: Harm, Probability and Another Look at Parfit's Depletion Example,” *Utilitas* (forthcoming 2006).

32. Gregory Kavka, “The Paradox of Future Individuals,” *Philosophy & Public Affairs* 11, no. 2 (1981): 93-112.

33. Just think: had your parents not attended just the schools that they did, would you have ever existed at all? Had they not married or fallen in love with each other just when they did, would you have ever existed at all? Had they had dinner at one restaurant rather than another the night you were conceived, would you have ever existed at all? The problem is not so much that new eggs are produced about every month or so. The problem is rather all those sperm: millions competing to inseminate a given egg during a given cycle, and many, many factors at play to determine just which one will succeed.

34. *Supra* note 31.

35. That we think probabilities have a role to play in determining whether harm has been imposed does not mean that *actual* wellbeing does not have a role to play in determining harm as well. Thus, we are free to say that an act that creates a less-than-maximal quantity of expected wellbeing for a given person *p*, but turns out, contrary to expectations, to be the objectively better-for-*p* act in light of how the future in fact unfolded, is an act that does not “harm” *p*. In the case at hand, however, this condition on harm would be satisfied given that the baby is in fact born damaged. Where this condition would come into play to avoid a finding of harm, then, would be the case where the babies born of ITISP “luck out” and are born healthy.