

Dan Moller

Abortion and Moral Risk

Suppose that you are considering performing some act A that you are worried might be wrong. Perhaps someone has presented you with an argument which purports to show that A is morally objectionable. Suppose, further, that after due deliberation you ultimately conclude that the argument fails, as do all the other anti-A arguments you know of. Is this the end of the matter? Should your deliberations be at an end after responsible consideration of the available arguments?

The surprising answer is “No,” according to a line of thought with roots in Catholic tradition and more recently echoed in certain decision-theoretical approaches to morality.¹ On this view, the mere *risk* of making a deep moral mistake rules out certain acts. If this were true, first-level deliberation about one’s actions would not be enough; we would need to proceed to second-level deliberation about the risk of being mistaken at the first-level, and doing so might rule out or at least count against A. To see the worry, suppose that you yourself believe that the argument counting against A is a powerful one, and that it only fails because of some intricate and subtle fallacy you have identified. In that case, there seem to be grounds to worry from within your own point of view: (1) it is quite easy to make mistakes about subtle philosophical arguments, and (2) if you have made a mistake, doing A would be seriously wrong. Going ahead with A under those circumstances makes it look as if you were willing to risk serious wrongdoing on the basis of beliefs about which you can’t (and perhaps couldn’t rationally be) very confident.

In the abstract, then, this line of argument may strike us as plausible or at least worthy of serious consideration, and when it is presented as a ground for something like vegetarianism, my experience is that philosophers are often willing at least to take it seriously. However, it can also raise disturbing questions about practices we may be more reluctant to question. In this paper, I attempt to build the general case for taking moral risk seriously, and consider how such an argument might apply to the case of abortion. Those of us inclined to support the legality of abortion will find such a line disturbing. And indeed, I don't think this argument can conclusively establish either that individuals must avoid abortion or that laws permitting abortion are wrong. But at least one version of the risk argument nevertheless seems to me strong enough to show that there is a moral *reason* to avoid abortion, albeit one that can be overridden in particular cases, and one that does not obviously support criminalizing abortion.

I – MORAL FALLIBILISM

Let me start out by briefly distinguishing the approach I intend to pursue from two others. Catholics have sometimes worried that we might be mistaken about what the fetus *is* (whether it is a human being, rather than a mere clump of cells, say). This approach is unpromising because there is nothing the fetus could be that would tell us straightaway that it was wrong to kill it. (Short of some question-begging moralized description, of course.) On some views, the fetus' turning out to be a human being wouldn't be enough for it to attain the status of psychologically complex persons who possess full-blown moral rights (Singer 1993, Tooley 1983). Moreover, there are writers

claiming to show that even if the fetus did have such rights it would still be permissible sometimes to kill it (Thomson 1971, Kamm 1992). So an argument from risk would need to be directed not simply toward the kind of thing the fetus might turn out to be, but also toward the moral principles involved. What we need is a more sophisticated approach—an argument that takes into account the possibility of error not just about what the fetus is but about the moral principles governing its treatment. The kind of fallibilism the argument from moral risk asks us to take seriously is fallibilism about our moral and perhaps related metaphysical beliefs.

On another view, having an abortion would be *irrational* on risk-related grounds (Lockhart 2000). Lockhart makes his case by asking us to assume that moral norms are the only norms that matter (p. 22), and then appealing to the analogy of expected utility theory (p. 27). The trouble, of course, is that in the typical case in which moral risk would arise, the woman will think that abortion is unlikely to be immoral, and have strong prudential reasons to abort. Telling her that although she finds the arguments against abortion unconvincing, she nevertheless has most reason from the *moral* point of view not to abort is unhelpful. The agent can reply, “Perhaps so, but those reasons seem weak, and my prudential reasons to abort are much stronger. So all things considered what I have most reason to do is abort.” Abortion thus seems unlikely to be irrational in any live case, at least for the sorts of reasons Lockhart cites.

The best way to tackle the issue of moral risk is rather to consider a two-stage argument that starts by showing that we are liable to making moral mistakes, and then goes on to identify norms applying to our susceptibility to such mistakes. In the rest of

this section, I try to establish the first point about liability to mistake, leaving the second point to section II.

By way of illustrating our susceptibility to error in the abortion case, consider a version of the Deprivation Argument.² That argument begins with the premise that killing someone is normally wrong because doing so deprives the victim of a future of value—of all of the future goods that the victim would otherwise have enjoyed. The second premise is that killing the fetus deprives it of a future of value, and that the goods the fetus is deprived of are comparable to those regular murder victims are deprived of. From this, the argument concludes, it follows that killing the fetus is deeply wrong.

Let us assume that this argument seems to fail—on balance, the objections strike us as more plausible than the argument itself. Why should we take seriously the suggestion that we might be mistaken in that conclusion? After all, there are many powerful objections that we might find convincing, and no doubt we have good reason not to begin doubting our own judgments without some strong impetus for doing so.

The problem is that the kinds of issues that are at stake in evaluating the Deprivation Argument are ones it is just very hard to be confident about. Consider two objections by way of illustration.³ First, the second premise assumes that the fetus has a future in which it enjoys certain goods, for instance the goods a young child experiences. That in turn presupposes that the fetus is identical to a future child, and this identity assumption may be denied. Proponents of psychologically oriented theories of personal identity may claim that the requisite psychological connections, the overlapping chains of memories, intentions, and so on, which license our identifying an individual existing at T with another existing at T', do not hold between the fetus and child.⁴ Although there may

be some differences depending on the details of the psychological theory in question and how developed the fetus is, on most theories the fetus is unlikely to meet the criteria for sharing sufficient psychological properties with a future child to count as identical with it, particularly in the first trimester when nearly 90% of abortions occur. In fact, we might even wish to put the point more strongly and say that the fetus is fundamentally an organism, whereas the child is fundamentally a psychological entity, and since these are distinct beings differing in their makeup, genesis and persistence conditions, it is simply false that the goods of childhood lie in the fetus' future.

However, there are rival accounts of personal identity that seem to support the second premise of the Deprivation Argument, particularly animalism (Olson 1997, van Inwagen 1990, Snowdon 1990). According to animalism, we are fundamentally creatures of flesh and blood—dynamic, homeostatic biological processes—and determinations of identity are to be made by investigating whether the two candidates are the same organism and share the same biological processes. According to animalism, it is not the case that some time after the creation of a human organism a distinct, psychologically constituted entity arises and goes on to enjoy various goods; rather, people are human organisms who came into being as soon as they met the basic conditions for the existence of an animal. Psychological properties are viewed as inessential to being and remaining the creatures we are.⁵ If this view is correct, then the fetus will probably count as the same thing as some future child and would therefore be deprived of future goods if it were killed. Animalism thus comports with the Deprivation Argument.⁶

Alternatively, we might attack the first premise for failing to incorporate the kinds of defeating conditions that Judith Thomson's defense of abortion brings out (1971). In

her famous case of the moribund violinist, an innocent person is forced by a third-party to render pregnancy-like aid to a total stranger. In response to the case, many people feel that the structure of rights and duties obtaining would make it permissible for the unwilling aid-giver to kill the innocent violinist, who, it is said, has no claim to the innocent's aid. Depriving the fetus of future goods can be defended, according to this line of argument, on the grounds that the structure of rights and permissions obtaining during a pregnancy permit a woman to kill the fetus. For the Deprivation Argument to be persuasive, we would need to believe that Thomson-style examples don't show that the fetus has no right to the continued support of the pregnant woman, and that abortion isn't tantamount simply to cutting off that support which isn't owed the fetus anyway (Kamm 1992).

But here, too, there are substantial grounds for doubt. The objection depends largely on the force of cases like the dying violinist, and numerous criticisms have been made of those cases and their relevance to abortion. One obvious worry is that they are applicable to cases of rape but not to more typical cases of non-coercive pregnancy. (If it turned out that the Deprivation Argument failed in the 1% of cases that involve rape, we could simply restrict its conclusion slightly and still produce a significant argument.) Critics, moreover, have pointed out various distorting features of Thomson-style examples. One of these emerges in recent work that seems to show that our reactions to cases are heavily influenced by whether we group victims together with their potential beneficiaries in a single situation or rather see them as inhabiting separate "moral zones" (Hershenov 2001, see also Unger 1996). When we are snatched away and forced to help someone avoid costs they face in their separate-seeming situation, we tend to think that

we are entitled to great latitude in avoiding the costs of aid. But if we think of ourselves as members of a group faced with costs arising in a common situation, we often think it wrong to avoid those costs by transferring them to our fellow group members. If this were right, then violin-type cases would begin to look questionable. They are usually designed so that the aid-giver is torn from her situation and then exploited by a third-party for her aid. This means that the cases seem to trade on grouping-reactions. If so, we face the problem of justifying the moral significance of our psychological tendency to group according to location, degree of mobility, and various other curious factors—otherwise we might just say that their force derives from morally irrelevant features.⁷ More importantly, the woman and the fetus actually seem to be in the *same* situation by the relevant criteria, implying that the analogies are crucially misleading in this respect.

So it looks like there is some non-negligible chance of being wrong about the Deprivation Argument. It is worth explicitly distinguishing the point being made here from others with which it might be confused. The claim is not that the mere existence of arguments against abortion or the mere fact that people disagree about abortion gives us a reason to reconsider our views.⁸ The claim is a narrower one concerning the possibility of error; nothing yet has been said about the doxastic or practical implications of that possibility. And even in that narrower context, the main reason for supposing there is a non-negligible possibility of error isn't the sheer existence of anti-abortion arguments. It is rather that the subject matter involved is the sort of thing it is all too easy for people like us to be mistaken about; abstruse moral reasoning involving far-out cases and complex principles is something we find very difficult and are disposed to get wrong

reasonably often. In this regard the claim about the possibility of moral error is no different than other kinds of errors.

Suppose that your husband has decided to bet the farm on a company's stock based on complicated calculations involving currency fluctuations, consumer demand, and so on. You point out arguments concluding that the stock will decline and suggest that he consider the possibility of error, but he insists that subtle economic analysis he has performed undermines those arguments and accordingly claims that any possibility of error is so low as to be negligible. ("The mere existence of counterarguments and the fact that many economists disagree shows nothing.") What is your husband missing? The obvious answer is that he is failing to consider how easy it is to go wrong when engaging in complicated financial reasoning. That may not be a reason for him to change his views, but it seems like a very good reason for him to acknowledge a significant possibility of error. Similarly, when doing complex moral reasoning, all of us are liable to make mistakes from time to time. Here, too, this may not itself give us a reason to change our views, but it does seem to be a reason to concede that there is a non-negligible possibility of error in assessing arguments like the Deprivation Argument, which is all that is being claimed at this stage.

Resistance to acknowledging the possibility of moral error may rest on how odd it can seem that we might know all of the facts involved in a case and yet be dramatically wrong about the morality. If there is not even the slightest uncertainty about the facts, are we really likely to be disastrously mistaken in our moral appraisal of the situation? The argument from risk requires us to believe that, despite our factual knowledge about the

fetus, the moral principles we embrace that cause us to reject the Deprivation Argument may be utterly mistaken. That may seem far-fetched.

But consider John Newton, author of “Amazing Grace,” slave-trader and (eventually) abolitionist. At some point, Newton came to believe that the people he had formerly sold into slavery had moral status and that his actions were morally repugnant. This change was neither a result of Newton’s religious beliefs (which he already held when he was captain of a slave ship⁹) nor of any new factual insight. Instead, Newton seems to have come to a realization that the same facts he was aware of earlier had a moral significance he had failed to register before: “The reader may perhaps wonder, as I now myself, that knowing the state of the vile traffic to be as I have here described, and abounding with enormities which I have not mentioned, I did not at the time start with horror at my own employment as an agent in promoting it. Custom, example, and interest had blinded my eyes” (Phipps 2001, 202). And when Newton proceeds to offer arguments against the slave trade, his arguments are distinctively moral in character. Although they often involve an emphasis on certain facts (specific ways in which slaves were treated) these facts are the same facts that Newton observed first-hand while still supporting slavery. “Surely, if the advocates for the slave trade [which has just been shown to involve rape] attempted to plead for it, before the wives and daughters of our happy land, or before those who have wives or daughters of their own, they must lose their cause” (Newton 1811, 53). Newton had long known of the rape-practices he refers to here; what is new is the moral insight. So Newton seems to be a clear example of how we can be grotesquely mistaken in our moral views even with complete knowledge of the non-moral situation. The same seems to be true of many other cases: the revisions in

people's attitudes toward animals or women or gays may often owe something to factual input, but other times the main stimulus is simply the realization that one's moral principles are awry.

A related worry is that appeals to the significant possibility of moral error presuppose moral realism. (“No moral mistakes without moral facts.”) But this, too, is off track. Many alternatives to moral realism offer surrogates for moral truth, and those surrogates yield surrogates for moral risk.¹⁰ This is almost inevitable since any alternative to moral realism will want to make room for the everyday experience of realizing we have made a moral mistake, and correlatively, for the experience of a moral conversion like Newton's. The concept of moral risk is rooted in pedestrian features of our moral lives that any theory will need to accommodate one way or another, setting aside radically revisionist theories that render normative ethics moot. To give one simple example, a theory might make sense of moral error by claiming that such errors (or “errors”) represent inconsistencies between our practices and beliefs, and our own deepest values. And notice that this wouldn't in any obvious way render the notion of moral risk trivial. Were most of us to learn that, unbeknownst to us, we were acting in a way that was cruel, degrading or wicked by our own standards, we would be deeply troubled. So different meta-ethical accounts will interpret talk of moral error and moral risk differently, but this should worry us no more than it should worry writers of garden-variety normative ethical theories that there are competing interpretations of what they mean by talking of moral truth.

II – NORMS FOR MORAL RISKS

Abortion involves taking a moral risk because there is a non-negligible possibility of wrongdoing even if we have decided that the arguments against abortion fail. And there is certainly a lot at stake: if we are mistaken and have an abortion, we are doing something deeply wrong. But that does not yet tell us what our attitude toward moral risks like abortion should be; that is a further normative question. We might, after all, agree that abortion presents a moral risk but decide to ignore it. This appears to be the view Simon May advocates in the arena of politics, for instance.

In the context of reasonable moral controversies, no person will be justified in adhering to her position without some measure of doubt, since her fellows may always raise relevant considerations—reasons of correction—that demonstrate that she has not, after all, attained true reflective equilibrium. But this doubt is normatively inert. Unless they actually do raise considerations that demonstrate this, she has no reason to think that she has not, in actual fact, formulated the best policy. That a political agent must admit the possibility of being wrong does not give her a reason to modify her position unless that doubt is accompanied by independent reasons that show she actually is wrong. (May 2005, 339-340)

If we are to take moral risk seriously we need an argument to show that there exists some moral norm governing our actions in light of such risks; we need to show that, contrary to May, the significant possibility of error is *not* normatively inert. We can approach such an argument by contemplating the various positions we might occupy in this area, which,

like ancient Gaul, is divided in three. These positions might be described as a moderate view flanked by two extremes. On the first extreme view, we never need to take moral risk into account; it is always permissible to take moral risks. On the second extreme view, whenever there is the slightest moral risk, we must refrain from acting; it is never permissible to take moral risks. The moderate position I will be defending is that we have a reason to avoid moral risk, variable in its strength, but not necessarily a decisive one, since it may be overridden by other considerations depending on the circumstances. So we should take all or nearly all moral risks seriously, but that does not mean avoiding any and all moral risks whatever the circumstances. Because the alternatives to this moderate view seem to be false, we have reason to take it seriously and consider the possibility that moral risk can give us a reason to avoid abortion.

To see why we should avoid the first extreme position that says we can always ignore moral risk, consider what that would imply. It would imply that even when there is a very significant chance of doing something catastrophically wrong, and even when avoiding that risk would cost us absolutely nothing, we may still ignore the possibility of wrongdoing. This implication is deeply counterintuitive. It is difficult to see how we could possibly care about morality while being unwilling to incur even a very small cost in order to avoid an enormous risk of doing something terribly wrong. If we are unwilling to incur such costs, we simply seem to be indifferent to moral norms; our stance vis-à-vis morality then seems little different from that of the amoralist.

To illustrate this, it is best to begin with a thinly described case. Suppose Frank is the dean of a large medical school. Because his work often involves ethical complications touching on issues like medical experimentation and intellectual property,

Frank has an ethical advisory committee consisting of 10 members that helps him make difficult decisions. One day Frank must decide whether to pursue important research for the company in one of two ways: plan A and plan B would both accomplish the necessary research, and seem to differ only to the trivial extent that plan A would involve slightly less paperwork for Frank. But then Frank consults the ethics committee, which tells him that although everyone on the committee is absolutely convinced that plan B is morally permissible, a significant minority—four of the members—feel that plan A is a moral catastrophe. So the majority of the committee thinks that the evidence favors believing that both plans are permissible, but a significant minority is confident that one of the plans would be a moral abomination, and there are practically no costs attached to avoiding that possibility. Let's assume that Frank himself cannot investigate the moral issues involved—doing so would involve neglecting his other responsibilities. Let's also assume that Frank generally trusts the members of the committee and has no special reason to disregard certain members' opinions. Suppose that Frank decides to go ahead with plan A, which creates slightly less paperwork for him, even though, as he acknowledges, there seems to be a pretty significant chance that enacting that plan will result in doing something very deeply wrong and he has a virtually cost-free alternative.

Is there anything to be said against Frank's decision? If so, the obvious candidate is that Frank is not taking what I have called moral risk seriously enough. That is, he is not placing enough weight on the mere possibility that his actions may involve himself (and the company) in wrongdoing. This seems to be true even though it is, on balance, somewhat more likely that enacting plan A is permissible. If this is how we judge the case, then we must reject the extreme position about moral risk and accept that

sometimes moral risk renders certain acts impermissible. However, the case is somewhat unusual in that the actual moral considerations that affect the case are screened off from the agent. We should consider a less thinly described case in which it is the agent himself deliberating over the relevant issues.

Suppose Sally has a temporary abnormality that will cause any child she conceives now to be severely handicapped, though not to the point that the child's life would not be worth living. If she waits a month, the child she conceives will be perfectly healthy (Parfit 1984, sec. IV). It strikes Sally that she has no moral reason not to conceive right now, since in doing so she will not be harming anyone. Her disabled child will have no complaint, since he cannot claim that he would have been better off without her action: assuming some form of genetic essentialism, had she not conceived when she did, a different egg and different sperm would have fused and the child would never have existed, so Sally's conceiving now, far from making the child worse off, actually causes him to come into existence. Having contemplated all this, Sally comes to the conclusion that conceiving now, if only because doing so would be a trifle more convenient for her, would be permissible. On the other hand, she has the lingering sense that perhaps there is something wrong with her act after all, even though she can't explain it in terms of how the act affects other people. Sally is aware that there are other ("impersonal") moral principles that could explain why it would be wrong to conceive now, but she judges that it is slightly more likely that those theories are mistaken. However, it also seems clear that if conceiving now is wrong at all it is probably not trivially wrong; the people who believe that there would be something wrong with conceiving now typically think that conceiving now would be *deeply* wrong. So the situation is this: on balance, Sally feels

that conceiving now is permissible, but she acknowledges that there is a very substantial risk that her act is deeply morally wrong. And she faces no significant costs for waiting to conceive; conceiving now would merely be a little more convenient. Nevertheless, Sally decides to conceive now and has a severely disabled child.

Now in considering this case it is very important not simply to think about whether what Sally does is right or wrong. For if we reject the principle that acts can only be wrong when they harm specific people, then it will seem obvious that what Sally does is wrong, but for reasons unrelated to risk: her act will be wrong simply because "impersonal" considerations, perhaps consequentialism, condemn it. What we need to consider is whether there is anything wrong with what Sally does *apart* from the obvious point that Sally's value theory may just be wrong. Or, to put it another way, are there grounds for criticizing conceiving now even from within Sally's point of view? It seems to me that there are. If nothing else, Sally should take seriously the possibility that in conceiving now she would be doing something deeply wrong; this should be of great concern to her, especially given that conceiving now, by hypothesis, does not do much to promote her interests. Ignoring this risk, as we must do if we are to claim that it does not provide us with at least some reason not to conceive now, seems to give insufficient weight to avoiding wrongdoing. This case, too, suggests we should avoid the extreme view that moral risk never matters.

The perspective of hindsight bolsters these considerations. We often look back and judge that we have been mistaken in some value judgment that we have made. Suppose you decide to have an affair because doing so would give you pleasure and you think that non-hedonic values are so much flimflam. After several years you may

abandon your hedonistic perspective and come to appreciate the force of values like loyalty and promise-keeping. Once you have attained this later perspective, you will look back with regret on the decision you made earlier, and that regret will be the result of your judgment that you made a moral (or axiological) mistake of some sort. Though these value-regrets are not commonplace, neither are they especially rare. (They are especially common in the prudential sphere: the wise old man looks back on his youth and regrets having attached such great value to making a fortune, and wishes he had spent more time with his family.) Now consider one of those acts that we regret at some point in the future because of the change in our values. Though we may or may not have known it, clearly there was some risk involved in acting as we did, since our regret proves that there was the possibility of a bad outcome. We may have lacked the moral sensitivity to register that possibility at the time, and so it could be a mistake to blame us for failing to take it into consideration, but our regret nonetheless does seem to show that the possibility was real and so too was the risk. Since we can envision the possibility of regretting our actions later due to our getting questions about value wrong, we seem to have a reason in the present to take seriously that possibility—not (merely) to avoid the pain of future regret, but to avoid the moral or axiological mistake that we might be making. Moral risk seems to be the forward-looking complement to the backward-looking regret we often feel about the value judgments we have made.

So cases like Frank's and Sally's, as well as the possibility of future regret, support the contention that the possibility of error about the moral principles relevant to a case gives rise to risk, and that risk gives us reason to refrain from acting in certain ways. There are norms governing moral risk, and sometimes these norms count against or even

rule out acting. And the worry, naturally, is that just as we may be mistaken in our appraisal of the significance of conceiving a child now or having an affair, we may be mistaken about the moral principles involved in the premises to the Deprivation Argument; if the uncertainty involved in the conception case gives rise to a pro tanto reason for avoiding conceiving now, so too uncertainty about the premises to the Deprivation Argument should give rise to a reason that counts against abortion.

If all moral risks were wrong, then conservatives would be right to infer directly that abortion is wrong, since the Deprivation Argument seems to show that we are taking a moral risk if we have an abortion. However, earlier I described the position that we should never take moral risks as an extreme one, and I do not believe it can be correct. This extreme view would imply that even when there is only a minuscule chance of wrongdoing, and even when avoiding the risk would involve staggeringly high costs to the agent, and even when, if the act were wrong, it would only be a case of trivial wrongdoing, we must nonetheless avoid the risk. These extreme implications are implausible. We are not required to ruin our lives in order to avoid an incredibly remote chance of trivial wrongdoing. (Readers can construct their own illustration around a charlatan purveying flimsy philosophical arguments for the view that we must incur some enormous cost in order to avoid a mere peccadillo.) And in any case, reasoning this way would seem to make permissible action impossible whenever competing moral considerations condemn both the act and its omission; in that situation, there simply is no avoiding moral risk. Presumably we don't want to say that an agent in that situation does wrong no matter what he does.

III – COMPLICATIONS

So far I have argued that (1) there is a non-negligible possibility of agents like Jessica being mistaken about the morality of abortion, and that (2) there are moral norms that count against taking such risks. However, I have also acknowledged that such *pro tanto* reasons to avoid moral risks might not be decisive, all things considered. How then do we determine which risks are permitted? My main goal in this paper is to explore the simple point that there is a reason to avoid abortion, often ignored, stemming from moral risk, and that that reason seems to be a non-trivial one. I am less concerned with working out the details of a complete theory of moral risk. But there are some obvious factors that agents like Jessica should take into account that are worth briefly discussing before returning to more fundamental problems with the notion of moral risk. (The list is not meant to be exhaustive.)

The examples already given make it obvious that moral risk is partly a function of

1. The likelihood that an act A involves wrongdoing, and
2. How wrong A would be if it were wrong,¹¹

and also that we should further adjust for

3. The costs the agent faces if she omits A.

All of the earlier cases rely heavily on stipulating the level of costs faced by the agent for not acting, the degree of wrongness that would be involved if the act were wrong, and the likelihood of wrongdoing. We can easily produce different responses to the kinds of examples we have discussed by adjusting these variables, which is evidence that they are important factors.

In addition, there are many other normative considerations that would enter into a full-fledged consideration of moral risk, besides the quasi-quantitative considerations mentioned so far. For instance, we would also need to consider

4. The agent's level of responsibility for facing the choice of doing A.

It is often clear that an agent is required to incur higher costs to avoid wrongdoing if the agent is responsible for facing the choice between wrongdoing and incurring costs in the first place. Agents who are simply thrown, by some third-party, into the position of having to make such a choice seem better able to justify rejecting costs and risking wrongdoing. And any thorough investigation of the norms governing moral risks would no doubt disclose additional features that are relevant to the moral acceptability of taking such risks, just as discussion of risks like drunk driving disclose subtle and complex factors that affect whether it is permissible to drive under various conditions.

One last factor should be mentioned—Lockhart's point from earlier that we don't face symmetrical moral risks in *not* having an abortion:

5. Whether *not* doing A would also involve moral risk.

Often, of course, rival arguments make differing prescriptions, and so there is risk either way. What we should consider, presumably, is the *net* moral risk involved in doing A—the risk involved in acting minus the risk involved in omitting. It is this issue that makes arguments from moral risk limited in application: such arguments will be most convincing when a large preponderance of the moral risks involved fall on one side or the other. When there are powerful moral arguments for both acting and omitting, then the argument from risk does not get a very strong grasp. Since there are few controversial issues where there are few or no moral risks on the one side, the argument from moral risk will be limited in its application. (Just to mention one additional example, though, consider vegetarianism. Animal rights proponents suggest that eating meat is deeply wrong since it allegedly contributes to massive unjustified suffering, while avoiding meat doesn't seem to be forbidden by any view. In this kind of case we can trade, as it were, moral risks for personal goods—"half the risk for half the taste"—and the arguments above suggest that sometimes we are *required* to make such exchanges. But there seem to be few moral quandaries that allow such exchanges.) Abortion, however, seems relatively unproblematic from this point of view, since there is usually little to be said for thinking that *having* an abortion might be morally required. By and large, pregnant women who decline to abort aren't doing anything immoral.

Lockhart seems to be right about this, but perhaps we should go into a little more detail. It is true, for instance, that there may be some situations in which it would seem wrong to gestate—perhaps when having more children would cause other community-members to starve, or when the child would suffer from some catastrophic birth-defect

that would render its life not worth living. In such exceptional cases we would have moral claims to consider both for and against having an abortion, and the argument from risk would be weak or silent. So much we should concede, though it is worth pointing out that to present a comparable risk, the conclusion to the argument for aborting would need to be quite strong. The Deprivation Argument claims that abortion is about as bad as killing an innocent person; to present a countervailing risk, the case against gestating would need to conclude that having a child was comparably wrong. But we might feel tempted to press further and point to the far more common situation in which young mothers have their lives shattered by the burdens of premature parenthood. This would be a mistake, however. Costs to the agent are clearly relevant to deliberations about risk, and were properly addressed earlier. But introducing them as countervailing moral risks would be inappropriate, since the costs to the woman don't create obligations (or even moral reasons) for her to abort. The burdens to the mother generate, at best, *options* for aborting, not requirements, for the same reason that someone has the option not to ruin his life to give his child a small pleasure, but is not morally required not to do so (Slote 1984). The person who ruins his life under those conditions might be crazy but he's not immoral. These considerations suggest that that the *net* moral risk involved in abortion is overwhelmingly on the side of having the abortion, at least in most circumstances. Because this is so, the argument from risk can get a hold on abortion despite factor 5 above.

Suppose that all of this were right. How strong a case against abortion would it yield? The main point is simply that agents should consider something like 1-5 when making abortion decisions; I will not try to show what exactly agents should conclude,

beyond the points made earlier—that it seems fairly easy to make a mistake about issues like personal identity, and that our obligation to avoid something as wrong as killing an innocent is presumably strong. In any case, as I have argued, the force of the risk-based point will vary with things like costs and facts about responsibility, which means that the force of this argument will vary from agent to agent. Clearly, then, the argument does not create a blanket all-things-considered objection to abortion; it does seem to suggest, however that there is a moral reason—probably not a weak one—for most agents to avoid abortion.

IV – CONCLUSION

I have described a general argument against abortion from moral risk, based on the point that the significant possibility of engaging in serious wrongdoing gives us a reason not to have an abortion. If this is correct, then we must not prematurely break off our moral deliberations once we have evaluated the soundness of the relevant arguments; we must go on to consider the risk that we are mistaken. Second-order reasoning has an important place in our decision-making. But I have also pointed out that in any particular case the morality of having an abortion will depend on various additional factors, such as costs to the woman. What has been said so far is compatible with abortion being quite defensible in situations in which the costs to the woman are high, perhaps because adoption is unavailable, or in which she bears little or no responsibility for facing the choice between absorbing the costs and taking a moral risk (e.g., rape). The implications of the argument from risk thus vary with the particular circumstances involved in any given pregnancy.

In order to show that in a particular case abortion is wrong we would need to construct a more specific argument than any I have given, one that would show that, given the particular circumstances, abortion would involve an impermissible risk. The argument from risk, then, shows that we generally have a powerful moral reason to avoid abortion (at least, if we grant that the Deprivation Argument is not hopeless), but it does not imply a blanket condemnation.

Throughout the paper, I have focused on the personal morality of abortion, not the public policy of abortion. It might be tempting to extend the reasoning above to the policy context, but we should be quite cautious about that. In the personal case, the agent sometimes has reason to accept higher personal costs (like foregoing an abortion) in order to avoid moral risks. But in the case of the legislator things look different: the legislator won't (typically) bear the costs imposed by criminalizing abortion—third-party women will. The structure of the legislator's decision is whether to impose large costs on many third-parties in order to avoid a moral risk. Suppose that abortion is permissible after all. In that case, imposing those costs would be deeply wrong. So in the case of the legislator, it is less clear that we can generate the kind of asymmetry in risk that we saw earlier is crucial to getting moral risk to apply. Obviously much more could be said about this, but in any case we should not assume that the extension of the personal case to be a straightforward matter.

Bibliography

- Blackburn, Simon (1998). *Ruling Passions*. Oxford: Oxford University Press.
- Cudd, Ann (1990). "Sensationalized Philosophy: A Reply to Marquis' 'Why Abortion is Immoral.'" *The Journal of Philosophy*, 87, 262-264.
- Gibbard, Alan (1990). *Wise Choices, Apt Feelings*. Cambridge: Harvard University Press.
- Hershenov, Scott (2001). "Abortions and Distortions: An Analysis of Morally Irrelevant Factors in Thomson's Violinist Thought Experiment." *Social Theory and Practice* 27: 129-149.
- Kamm, FM (1992). *Creation and Abortion*. Oxford: Oxford University Press.
- Kelly, Thomas (2006). "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology*, John Hawthorne and Tamar Gendler Szabo (eds.). Oxford: Oxford University Press.
- Lockhart, Ted (2000). *Moral Uncertainty and Its Consequences*. Oxford: Oxford University Press.
- Marquis, Don (1999). "Why Abortion Is Immoral." In *Bioethics*, edited by Helga Kuhse and Peter Singer. Oxford: Blackwell.
- (1989). "Why Abortion is Immoral." *The Journal of Philosophy*, 86, 183-202.
- May, Simon (2005). "Principled Compromise and the Abortion Controversy." *Philosophy and Public Affairs* 33: 317-348.
- McInerney, Peter (1990). "Does a Fetus Already have a Future-Like-Ours?" *The Journal of Philosophy*, 87, 264-268.
- McMahan, Jeff (2002). *The Ethics of Killing*. Oxford: Oxford University Press.
- Newton, James (1811). *Works*, VI.
- Norcross, Alastair (1990). "Killing, Abortion, and Contraception: A Reply to Marquis." *The Journal of Philosophy* 87(5), 268-277.
- Olson, Eric (1997). *The Human Animal*. Oxford: Oxford University Press.
- Phipp, William (2001). *Amazing Grace in John Newton*. Macon: Mercer University Press.

- Pfeiffer, Raymond (1985). "Abortion Policy and Uncertainty." *Social Theory and Practice* 11: 371-386
- Singer, Peter (1993). *Practical Ethics*. Cambridge: Cambridge University Press.
- Sinnot-Armstrong, Walter (1999). "You Can't Lose What You Ain't Never Had: A Reply to Marquis on Abortion." *Philosophical Studies*, 96, 59-72.
- Slote, Michael (1984). "Morality and the Self-Other Asymmetry." *The Journal of Philosophy* 81: 179-192.
- Thomson, Judith Jarvis (1971). "A Defense of Abortion." Reprinted in *The Rights and Wrongs of Abortion*, M. Cohen, T. Nagel, T. Scanlon, eds. Princeton: Princeton University Press, 1974.
- Tooley, Michael (1983). *Abortion and Infanticide*. Oxford: Oxford University Press.
- Unger, Peter (1996). *Living High and Letting Die*. Oxford: Oxford University Press.
- Van Inwagen, Peter (1990). *Material Beings*. Ithaca: Cornell University Press.

¹ See, e.g., §13, “Declaration on Procured Abortion,” Sacred Congregation for the Doctrine of the Faith 1974, and the discussion of Lockhart below. There has been some previous discussion of the issue in Pfeiffer 1985.

² This is (roughly) the Deprivation Argument advanced by Don Marquis 1989, 1999. It is important to point out that to keep the discussion manageable, I am focusing on a single particularly interesting objection to abortion. A full discussion would need to consider *all* such arguments and the risk that *any* of them goes through.

³ Just as there are many arguments against abortion, there are many objections to the Deprivation Argument besides the two I consider here. See, e.g., Cudd 1990, Norcross 1990 and Sinnott-Armstrong 1999.

⁴ See Parfit 1984 for an account (not an endorsement) of the psychological theory of identity, as well as further references. For a more recent psychologically-oriented theory, see McMahan 2002. McMahan makes an objection to Marquis closely related to the one I am describing, as does McInerney 1990.

⁵ Van Inwagen: “The capacity for thought...seems to be, metaphysically speaking, a rather superficial property of myself” 1990, 120-121.

⁶ But notice that even on animalism very early fetuses may not count as identical with future children if it turns out that there is a stage when the fetus is not yet a biological organism. Perhaps the best account of organisms will fail to categorize the zygote or pre-embryo as an organism, especially when twinning is still possible. (This would probably affect fetuses less than two weeks old, at which point twinning is no longer possible.) For more discussion, see van Inwagen 1990, 152-153.

⁷ See Unger 1996, 101-106 for some disturbing analysis.

⁸ Several writers have given reasons to doubt this claim, among them Kelly 2006 and, in a public policy context, May 2005.

⁹ A biographer comments on Newton's post-conversion diary: "His diary contains dozens of pages largely filled with his confessions of personal wrongdoings, but one searches in vain for any awareness of sin in relation to his slaving business. If one did not know what his ship's mission was, one might presume that he had been on an extended spiritual cruise for meditating about his Savior. By way of declaring that his work was a religious enterprise, Newton began the ship's log for each voyage with the invocation, *Laus Deo* (by God's permission [sic]), and concluded it with *Soli Deo Gloria* (solely to God's glory)" Phipps 2001, 62.

¹⁰ I'm thinking of familiar expressivist stories such as Blackburn 1998 and Gibbard 1990.

¹¹ Factors like 1 and 2 mustn't be interpreted with an artificial degree of precision. It is unclear, for instance, that the relevant functions are straightforwardly linear. It is also unclear how exactly to interpret probabilities in this context. However, these and other issues that would afflict attempts to formalize the argument here seem to pose a similar problem for homely cases like the risky investment.