

Applied Moral Philosophy

Richard Arneson

7187

Applied moral philosophy is the application of moral philosophy to the issues, what laws and public policies should be enacted and implemented, and what individual actions ought individual agents to choose and execute. More broadly, applied moral philosophy might encompass the assessment of particular social practices and social arrangements, social norms, and actual moral doctrines prevailing in particular societies at particular times, according to the normative standards supplied by fundamental moral principles.

So described, “applied moral philosophy” might sound pleonastic. What on Earth would nonapplied moral philosophy amount to? Isn’t guidance of conduct and assessment of practices the point of moral thinking?

This is not a purely rhetorical question. One could take a purely intellectual interest in moral philosophy questions. At least, such an enterprise might make sense on the assumption that moral realism is correct. Some hold that moral claims are not genuine assertions but are rather more like expressions of attitudes or maybe commands. But according to moral realism, moral claims are genuine assertions, which admit of being true or false, and some moral claims are true. (That is to say, moral claims are not one and all false in virtue of embodying the false presupposition that there are moral facts to be discovered, which are potential truthmakers of moral claims.) Suppose moral claims so understood concern what we owe to each other—what is the due concern-and-consideration that any rational agent owes to any rational agent and to which any rational agent is entitled to get from others. An intelligent predator from another planet, who has no intention whatsoever of ever conforming its conduct to any moral demands from concern for what reasons there are to give due concern and consideration to others, might yet have an intellectual curiosity as to what the content of this due concern and consideration is. The curious predator being imagined has no desire to apply the results of this inquiry to any recommendations for anyone’s conduct or policy.

But this still seems a distinction without a difference. If our curious predator produces advances in moral philosophy, these if broadcast would be available as useful guidance to anyone who wants to do what is morally right and to advocate for moral improvement in laws, public policies, social practices, currently entrenched conceptions of morality to which people have allegiance, and so on.

A more fruitful way of characterizing the field of applied moral philosophy would begin by noting that questions about what is proper and right and what aims are worth pursuing can be asked at very different levels of abstraction. One can ask moral questions about what a particular individual did at a certain time in particular circumstances. Was Sally right to fire her employee Al who was just emerging from a troubled marriage and showing every sign of renewed dedication to his work responsibilities? Was Ali behaving impermissibly in spreading gossip about his friend Jose, given the risk that the rumors would wreck his relations with some family members and disrupt his relations with almost all of them? One can engage in moral criticism of moral rules that are upheld by one or another society, or by some subset of the members of a particular society, the rules being implicitly or explicitly regarded as governing what to do only given that certain background facts obtain. For that matter one can engage in

moral criticism of systems of rules and social practices that prevail in particular societies at particular times, whether or not the rules in question are deemed to be moral rules or rules of some other sort. One might also attempt to formulate general moral principles for the regulation of conduct, the principles being regarded as norms that apply not just to one society at a time but across societies and times. “General principles” in turn might be more or less general. Candidate moral principles might aim to be authoritative for all humans in modern societies, for all humans, for all rational agents given general background empirical conditions that frame all action and choice, or for all rational agents or partly rational agents in any metaphysically possible or (more encompassing yet) all logically possible circumstances.

Against this backdrop, we can understand applied moral philosophy to be a designation for moral philosophical inquiry directed at more concrete and less abstract levels. This is vague, but what reasonably counts as applied moral philosophy is perhaps fixed only vaguely and in ways that are context-dependent. Inquiry directed at the level of fundamental moral principles—norms that if true at all are supposed to be true universally, necessarily, and timelessly--does not qualify as applied moral philosophy.

In the same spirit, we might alternatively identify applied moral philosophy as inquiry directed at moral concerns as they arise in everyday decision contexts and institutional design contexts and are identified as puzzles and problems by nonphilosophers seeking to go about their lives and conduct their affairs.

1.

Terms like “applied moral philosophy” or “applied ethics” might suggest an alternative picture of the enterprise than the idea of ethics directed at practical problems of personal and social life. We might conceive of applied ethics on analogy with applied mathematics or applied physics. In both cases a known body of knowledge in the form of general principles is applied to solve some less theoretical, more problems. (Or at least, this is our simple and naïve idea. For introduction to some real-world complexities of applied mathematics, see the discussion in Wilson (2006, chapter 8) of the nineteenth-century engineer Oliver Heaviside’s innovative treatment of linear differential equations to solve such problems as sending undistorted electrical signals over long distances.) On this model, to practice applied ethics one needs to have on hand a set of moral principles known to hold universally, necessarily, and certainly. For the set of principles to give determinate practical guidance, either the set must consist of one master principle, or include a determinate ordering of priority if there are several, which amounts to the same thing. With the moral principles specified and their meaning made clear, the application of the principles involves discovering and articulating the actual empirical facts of our situation that (according to the principles) determine what act to do (or policy to adopt) in that situation.

The fly in the ointment is that there is rampant disagreement about what are the correct moral principles or even if there are any such things. Different societies at different times embrace different moral codes, and within any given society at a given time there is a lot of disagreement among persons as to what general moral rules or principles are true. In this predicament, a reasonable person might have some hunches about where general moral truth lies, but will lack full confidence in his hunches. He or she might be wrong, and some others right. Or maybe the moral truth that determines

what to do in a domain that concerns us has not been discovered. If we do not know what the true determinate set of moral principles is, that could form the basis of applied ethics as described a few sentences back, then we cannot practice this top-down method of applied ethics. Moral disagreement is rampant not only among ordinary citizens and especially among people who have not thought much about these matters. Among people who have reflected hard on the issues and among academic philosophers who specialize in the study of ethics, moral disagreement is deep and pervasive.

If we cannot practice a top-down method of using known general principles to derive practical judgments of what to do in particular circumstances, will a bottom-up method suffice? This would involve formulating intuitive responses to particular questions such as “Is it morally required that Susan right now abandon her relationship with the somewhat abusive Fred?” and “Was it wrong for the U.S. to drop atomic bombs on Hiroshima and Nagasaki toward the close of World War II?”. It is hard to find moral magic in reliance on particular judgments. First, disagreement about particular judgments is about as pervasive as disagreement about general claims. In the face of disagreement, how can one gain confidence that one’s own view is any more likely than any other to be correct?

Second, anyone’s current moral beliefs depend on socialization and experience, which may induce bias or other suspect processes of belief formation. The attempt to put one’s moral beliefs in systematic order does not necessarily correct the problem, but may help.

A possible way to proceed is to combine top-down and bottom-up approaches, reflecting on one’s current beliefs and seeking to reflect on them critically and revise them with a view to arriving at an equilibrium state of belief in which one’s particular judgments and general judgments are mutually consistent and hang together coherently. A more general claim may merit rejection because in specified possible circumstances it implies recommendations as to what to do in specified particular circumstances that one finds unacceptable. A more particular moral claim may merit rejection because it conflicts with general rules or principles that explain and justify many of one’s beliefs and that seem plausible after scrutiny. We go back and forth. This zig-zag procedure has been called the “reflective equilibrium” method (Rawls 1999, chapter 1).

According to the reflective equilibrium procedure, in our thinking about what morality requires and permits we seek to find an ideally coherent complete set of ethical judgments. A complete set gives determinate advice in any situation one might conceivably encounter without need for any further evaluative judgment. Given the principles already on hand, learning the empirical facts that those principles themselves single out as relevant for choice suffices to determine correct choice.

But nothing in the reflective equilibrium procedure guarantees that we can attain knowledge of such a complete set of judgments. Nor is it guaranteed metaphysically that such a set must exist, whether or not our moral perception and moral reflection abilities would enable us to find it.

An individual starts with a set of moral intuitions. These may be at different levels of generality, down to specific judgments about what particular persons should do in particular circumstances at a particular time. A moral intuition is a judgment about how things seem to the individual doing the judging. It seems that Fred should repay the debt, in these circumstances. It seems that slavery is wrong.

One's starting point set of moral intuitions is very likely to be at least latently inconsistent. One might believe that all humans are equally morally considerable, that possession of greater rational agency capacity makes one more morally considerable, and that humans differ in their possession of rational agency capacity. Not all three claims can be true; something has to give. Some things one believes yield implications for what is morally right that conflict with the implications regarding what is morally right that are delivered by other things one believes equally strongly, at least in some conceivable circumstances.

Even when not inconsistent, one's starting point beliefs can seem a jumble, a pile of disparate independent claims. Particular moral claims are implicitly general claims. In making a particular claim one at least implies that the same judgment should hold in relevantly similar circumstances and not all differences in situations amount to relevant dissimilarity in circumstances. A particular claim is explained and justified by a moral general claim from which it can be derived; and the same goes for limited general claims that are hemmed in by conditions.

Describing reflective equilibrium method, John Rawls (1999) observed that in moral thinking one seeks a state of reflective equilibrium in which the particular claims one accepts after reflection and scrutiny are explained and justified by general claims one accepts after reflection and scrutiny. One's system of beliefs, for now, coheres. But further reflection and argument can throw this set of beliefs into disequilibrium again. One then seeks to restore, or at least move toward, coherent harmony. We can stipulate that ideal reflective equilibrium is the set of moral judgments at all levels of generality that an ideal moral reasoner who thinks clearly and properly evaluates arguments and integrates the results of her reasonings into a coherent set would arrive at after ideally extended ethical reflection that encompasses all relevant arguments in light of all relevant empirical facts. (A further claim would be to identify this epistemically ideal state with ethical truth: the moral facts are whatever would be believed in ideal reflective equilibrium.)

Described this way, the reflective equilibrium idea is not one candidate among others for how to characterize moral thinking and moral progress. As T. M. Scanlon (2003) has observed, it's the only game in town. This means it does not settle in advance any ethical controversies. It could be the case that there are no true general moral principles; moral truths are all particular judgments. It could be the case that there are no moral truths at all and sustained thinking straight about ethics would lead to skepticism. Or perhaps foundationalism in ethics is correct: a rationally defensible moral claim must either be or be derivable from a foundational claim that is self-evidently true. If foundationalism is correct, then as we approach ideal reflective equilibrium, the set of ethical beliefs we accept takes the shape of a foundationalist system with self-evident moral propositions at its bottom.

Acceptance of reflective equilibrium methods of ways of seeking moral truths might seem to assign a very important role to applied moral philosophy. Any moral theory in the sense of a proposed set of fundamental principles, however elegant and grand and even sublime it might seem, can be refuted by the counterexample of a single very plausible particular judgment that contradicts the theory! But equally, our best answer to any practical moral controversy that confronts us and to any perplexity about what it is morally right to do in very particular circumstances might require formulation

and examination and defense of an abstract general moral principle! On this way of understanding the moral enterprise, no clear or for that matter important line separates applied ethics from ethics proper. There is just moral thinking, which seeks to be both particular and general. At most, there is the distinction between moral thinking that is directed toward solving practical, real-world issues about what to do and moral thinking that is not directly concerned about these questions.

2.

The preceding discussion does not take us very far. According to the reflective equilibrium idea, one should try to arrive at moral judgments concerning real-world moral controversies by intuitively balancing one's particular and general prereflective judgments that bear on the controversy at hand, reflecting critically on one's resulting views, and continuing the process. Let us assume this is right as far as it goes. But where does such a process take us? How do we actually proceed, and to what judgments does reflection take us?

The short answer is that, so far, different individuals undertaking this process arrive at opposed beliefs. To illustrate, let us consider some opposed views that prominent theorists doing what is recognized as applied ethics have defended. These vary along several dimensions.

2a.

In the "Preface" to his book *Moral Thinking* (1981), R. M. Hare suggests that the stakes involved in trying to get straight about moral theory questions are extremely high:

"I offer this book to the public now rather than later, not because I think it needs no improvement, but because of a sense of urgency—a feeling that if these ideas were understood, philosophers might do more to help resolve important practical issues. These are issues over which people are prepared to fight and kill one another; and it may be that unless some way is found of taking about them rationally and with hope of agreement, violence will finally engulf the world."

What Hare has in mind is that people do not understand the nature of moral prescriptions, the requirement of universalizability, and hence the logical inescapability of acceptance of preference satisfaction utilitarianism, at least if one understands what one is saying and is thinking straight. Moral judgments are prescriptions constrained by universalizability. Making a moral prescription, one must be willing to prescribe that what one is urging some to do here and now should be done by anyone anywhere in comparable circumstances, including by oneself if one found oneself in such circumstances. Each of us seeks to satisfy his own preferences, and universalized, this becomes the prescription that overall preference satisfaction by anyone anywhere be maximized.

Hare is making at least two separate and distinct large claims. One is that moral judgments are really prescriptions, not assertions that could be true or false. The second claim is that once we understand the basic moral concepts, we must accept that the only sensible moral prescriptions are those that conform to preference satisfaction utilitarianism—we must prescribe that what should be done is whatever would maximally satisfy the preferences of all people who might be affected by what is done. Only confusion stands in the way of moral agreement.

The first claim is about what it is to make a moral judgment; its bearing on what to do is uncertain and possibly nil. Set that aside. The second claim if correct would be of immense practical importance, as Hare notes. But hardly anyone agrees with Hare that if you reject preference satisfaction utilitarianism you must somehow be logically confused.

2b.

Hare's conjecture that gaining agreement at the most abstract level of fundamental moral principle and metaethical understanding of the nature of moral claims and argument is necessary for philosophers to contribute to practical moral issues is contestable. An opposed view is that one can start in the middle, as it were. One looks to such principles and moral considerations as one can find that promise to have a bearing on the question, what resolution is the practical issue under discussion is most reasonable. Joel Feinberg defends the strategy of starting in the middle and eschewing grand moral theories. In the "General Introduction" to his four volume treatise on *The Moral Limits of the Criminal Law*, in which he defends the claim that only harm to others and offense to others justify attaching criminal law punishments to a type of conduct, he writes that these seeking such grand theory will find in his work "no semblance of a complete moral system, no reduction of moral derivatives to moral primitives. . . . It would be folly to speculate whether the moral theory implicit in this work is utilitarian, Kantian, Rawlsian, or whatever. I appeal at various places, quite unselfconsciously, to all the kinds of reasons normally produced in practical discourse."

The potential trouble with this mid-level and eclectic approach is that it occludes relevant possibilities. One is that some initial moral judgments about particular cases or middle-level principles would be exhibited as anomalous outliers in the light of a candidate universal principle that turns out to have plausible implications across the entire spectrum of possible applications. On this ground the outlier judgments should perhaps be discarded. This enlightenment only arises if we are willing to examine the most general principles.

This is potential trouble, but Feinberg will respond that in the four volumes of his treatise he carefully analyzes the ideas of legal coercion, criminal law prohibition, harm to self, harm to others, offense to others, and harmless wrongdoing. He surveys a wealth of legal and other examples and suggests that our considered response to these examples should be to embrace a general claim about legal paternalism: restriction of someone's liberty to act as she chooses, imposed against her will and for her own good. The general claim is that the use of the criminal law to restrict people's substantially voluntary choices on the ground that they are self-harming is never morally acceptable. The considerations that support this judgment, according to Feinberg, probably support a broader judgment to the effect that no restriction of a person's substantially voluntary choices against her will for her own good, whether by criminal law prohibition or any other means, is never morally acceptable. Examination and reflection on many types of case in which overruling people's substantially voluntary choices for their own good is either perpetrated or proposed in modern societies lead us to embrace the general claim just stated, which amounts to the view that each person is sovereign over her own life, has the right to live as she chooses, even to go to ruin in her own preferred way, so long as she does not thereby wrongfully cause harm or offense to others. So Feinberg argues.

2c.

Feinberg's sweeping rejection of paternalism stops short of rejection of enforceable requirements that people extend aid to others, help them when they are in peril, for the sake of the imperiled. Restriction of liberty to prevent harm to others can be morally acceptable according to Feinberg. Our duties to improve the world by helping to prevent harm to others set a limit to personal sovereignty.

If we were to accept this line of thought, the next question that would arise is, under what conditions do we have moral duties, maybe enforceable moral duties, to act at cost to ourselves to prevent harm to others that we ourselves have not been responsible for bringing about in the first place. Peter Singer advances a provocative answer to this question in a famous essay (1972). He appeals to a single example, invites our judgment about that example, and suggests a general principle that explains and justifies the judgment about the example that he thinks has common sense support.

Here is the example: Suppose you were walking past an isolated pond, and noticed a child drowning in shallow water. Surely any decent person would feel duty-bound to save the child. Singer suggests this verdict has immediate implications for the question, what if anything do we owe to distant needy strangers in peril.

Singer proposes this principle: "If it is in our power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it." The phrase "without sacrificing anything of comparable moral importance" blocks the duty from forming when one could prevent something bad from happening only by doing something wrong in itself, or by ignoring some alternative act that would prevent worse harm, or by also causing some comparable harm. The principle explains and justifies the judgment that one should save the child drowning in the pond, but also has much wider reach. If one can prevent harm to distant strangers by contributing to a disaster relief fund, the principle says one should give, and keep giving, until further giving would bring about a loss to oneself as great as the harm one could prevent by the next act of giving. According to Singer's principle, what we ordinarily regard as morally optional charity is in a wide range of actual circumstances morally required harm prevention.

Singer's way of arguing appeals to the thought that any judgment about what to do in particular circumstances must be derivable from some general rule one accepts. This is a requirement of articulate consistency. If one accepts that one ought to save the child in the drowning pond scenario, there is pressure to show that this claim is explained and justified by a general rule. One might resist the thought. But one might accept it while offering principles that comport with a principle of easy rescue but stop short of affirming stronger duties to help the needy.

2d.

In a brilliant essay in practical ethics, Judith Thomson (1971) deploys common views about easy rescue and the limits of the moral requirement to save needy persons from peril to unsettle debates about the moral permissibility of abortion. Many discussions of the intentional termination of pregnancy assume that the issue turns on the moral status of the fetus. Is the fetus morally considerable, a being whose interests have some weight in determinations of what to do, or not? If the fetus is morally considerable, does it have the full moral standing of personhood or some lesser status? Many hold that

if the fetus is a person, then abortion is morally wrong and should be prohibited by law. The background assumptions here seem to be that abortion is killing the fetus, that a being with the full moral status of personhood has a right to life, and that killing a being with the full moral status of personhood would violate its right to life, and thus would be morally wrong.

Not so fast, urges Thomson. Abortion withdraws the mother's body from the fetus, who needs it, but to show this withdrawal of service would be morally wrong, one needs to argue that the fetus has a moral right to the continued use of the pregnant woman's body. According to Thomson, it would be charitable for the woman to give over her body for the fetus's use through nine months of pregnancy, but as a general rule charity of this degree of magnitude is not morally required, and coercing people to be charitable to this degree of magnitude would be morally outrageous.

To illustrate the point, she has us imagine that a man finds himself attached to an ailing violinist, with tubes circulating nutrients from his bloodstream to the violinist, who is sick, and will die if this connection, set in place by his friends, is not maintained. Thomson's intuition is that it is plain as day in this scenario that the man has no obligation to provide needed care to the ailing violinist, and is morally at liberty to sever the connection, even if this withdrawal of aid, to which the violinist has no moral claim, will cause his death. Thomson comments that in law and morality we do not require even minimal samaritanism, so it is an enormous anomaly, involving discrimination against women, if law and morality require pregnant women to be splendid samaritans toward fetuses that come to be living in their bodies and dependent for life upon their willingness to go sustain the pregnancy for nine months until birth. Even if we should decide that some minimally decent samaritanism when some find themselves uniquely positioned to offer lifesaving aid should be required, this would still leave unjustified the treatment of women in a society that prohibits abortions across the board.

As many commentators have noticed, Thomson's example succeeds both in raising the issue, what level of charitable giving toward those in need, if any, should be required by law and morality and coercively enforced, and also raises the issue, whether blanket prohibition of abortion is morally unfair to women.

There are complications to be noted, and Thomson's essay flags them. In the ailing violinist example, the unwilling helper is connected to the person who needs help without the former's agreement and without any act or omission on his part that could reasonably be thought to have triggered a duty of care toward the ailing violinist. In the situation of pregnancy, matters are different. If the pregnancy arises from consensual sexual relations, then the woman who is now pregnant engaged in sex knowing that a pregnancy might result, and that in that case a person would be dependent for its life on her continuing aid during the term of pregnancy. Or at least, the woman ought to have known this fact, and is perhaps culpably ignorant if she does not.

Thomson observes that in the case of pregnancy arising from rape, the analogy to the original ailing violinist example is quite close. Other cases will vary along various dimensions, and in these cases, a verdict that the woman has acquired a duty of care toward the fetus that justifies requiring her to continue the pregnancy is not so clear. Perhaps the woman took prudent steps to avoid pregnancy, and she had bad luck. At any rate, in the case of pregnancy, the child would not exist but for the pregnancy, and is

arguably no worse off for existing and being given aid for a short period, even if the woman then decides not to continue the pregnancy and withdraws aid (Kamm 1992).

The upshot, as Thomson sees it, is that she has shown that the moral right to life, thought to be simple and transparent, is in fact complex and opaque. The right not to be killed is more accurately characterized as the right not to be wrongfully killed, so the plain fact that abortion standardly kills the unborn fetus does not settle the question, whether abortion violates any moral right possessed by the fetus. So even if, for the sake of the argument, we assume what Thomson believes to be surely false, that the fetus is a person with all the moral rights of a person from the first moment of conception, the moral status of abortion remains delicate, and there are strong arguments in support of the permissibility of abortion.

Thomson's reasoning is powerful, but can be run in the reverse direction. If we believe, with Singer and others, that moral duties to aid the needy are generally stronger and more far-reaching than common opinion tends to allow, then if the fetus is a person with full moral standing, its helpless dependency during pregnancy surely triggers strict moral duties to aid (though perhaps the cost of aiding should be widely distributed among all of us, not left to fall solely on the pregnant woman), and the prospects for defending the moral permissibility of abortion become dim (Warren 1973, also Marquis 1989, but see also Kamm 1992 and Boonin 2003). The focus of inquiry then turns back to the moral status of the fetus.

Here's one view. So long as the human fetus lacks sentience, it is not morally considerable at all, and is not an object of moral concern. Even after acquiring sentience, and becoming morally considerable, the human fetus lacks rational capacities necessary for personhood, so is morally considerable only at the level of a nonhuman animal that lacks rational capacity, such as a chipmunk, and lacks a serious right to life of a sort that would render abortion impermissible or seriously problematic.

This position seems to accord with common sense, and to be persuasive to the secular mind, until one notices that not only the young fetus, but also the late-term fetus and even the newborn infant, and beyond that the young child, also appear to lack the developed rational capacity that suffices for personhood. So the position just described seems to be on strong ground in robustly supporting the moral permissibility of abortion but to be in the swamp, or at least perplexing, insofar as it also robustly supports the moral permissibility of infanticide.

The enduring greatness of Thomson's essay lies in its clear delineation of the moral basis of the view that abortion consists in denial of lifesaving aid to the fetus that the pregnant woman is not necessarily morally bound to provide. In this delineation the abortion issue is revealed to be tied to fundamental issues concerning the content of the basic moral principles that specify what we owe to one another. There are several such issues, and they interact in complicated and contestable ways. The essay contributes to our understanding of a practical moral controversy by pressing us to accept that an acceptable position on a topic such as abortion cannot be anomalous but must fit within a coherent set of general principles the implications of which for other questions and other types of controversy are acceptable after reflection and scrutiny.

Thomson's use of analogies and philosophical thought-experiments also raises deep questions about moral theorizing. She is discussing a controversy about what we ought to do and allow in actual circumstances we presently face. She presses her point of

view by introducing unrealistic counterfactual scenarios, and asking what we ought to do in settings we certainly will never face. The background assumption is that an acceptable moral principle must yield implications for conduct not only in present and likely scenarios and in possible scenarios close to the actual world but also in remote possible world situations that are logically possible but physically impossible. Is this assumption acceptable? Opinions differ.

2e.

The question, what traits must being possess to qualify as a person, with the rights and duties of persons, is flanked by another: what makes a being morally considerable at all, so that its interests affect what it is morally permissible and morally required to do, when what one might do might affect the being's interests? The questions introduce the large issue, what morality requires of human persons, in dealing with nonhuman animals. On some views, animals that are not persons are just there for the use of persons and are owed nothing. On some views, the interests of nonperson animals count, but at a discount, when their interests conflict with the interests of persons. From a third standpoint, all animals are equal, in the sense that if, for example, a young rat or a young human would get the same pleasure from eating a morsel of food, the pleasure either would get should have identical weight in the determination of what is morally permissible and morally required to do.

2f.

Suppose we accept the idea that whatever morality demands of us, it definitely includes a strong beneficence component. That is to say, a large component of the demands that morality imposes on us consist of a general requirement to bring about better outcomes, with outcomes assessed by impartial standards. We might then think, in for a penny, in for a pound. Maybe the demand of morality is beneficence, period.

So construed, morality requires that one ought always to do whatever would lead to an outcome no worse (according to impartial assessment) than the outcome of anything else one might instead have done. This doctrine is known as act consequentialism (on varieties of consequentialism, see Sinnott-Armstrong 2011). It conflicts with many people's moral beliefs, by denying both moral options and moral constraints. One has a moral option when there is course of action one is permitted to do, that would not bring about the best consequences one could obtain.

Moral constraints are limits on the treatment one accords to others affected by what one might do. They comprise familiar moral prohibitions—Don't lie, Don't cheat or steal, Don't harm innocent nonthreatening persons, Don't break your promises or contracts, Don't force or coerce others behaving innocently, and so on. Such constraints might be regarded as absolutes, never to be violated. Most people probably regard them as serious considerations, but not exceptionless (Ross 1930). They might themselves come in conflict as when one can keep a promise only by telling a lie. They might be viewed as giving way when the impartially assessed consequences of conforming to them would be "sufficiently" bad. One might conceive of moral constraints as moral rights—claims possessed by a person that consist of duties on the part of specified others owed to her.

In many settings, violating moral constraints would also bring about less good consequences than conforming to them. But when one could bring about a slightly better outcome overall by oneself violating some moral constraint, there is rock-bottom moral disagreement between act consequentialists and proponents of any version of a morality of constraints or individual moral rights.

The disagreement between act consequentialism and constraint-based (or rights-based) morality looks to be a disagreement at a high level of moral abstraction, but it also looks to have immediate practical implications regarding what is the thing to do in many concrete real-world decision problems. Of course, where exactly the points of disagreement lie depends on what standards one ought to employ for the impartial assessment of outcomes/consequences and on what exactly the content of people's individual moral rights constraining how one may treat others is determined to be. In what circumstances is it wrong to engage in paternalism—restricting someone's liberty against her will for her own good? This depends on whether there is a moral constraint against paternalistic interference along the lines Feinberg and others propose. From a consequentialist standpoint, the issue turns on what the consequences of paternalism and nonpaternalistic alternatives would be and on the right way to assess these consequences. If a person's life goes better, other things equal, when she is autonomous or self-directing, the value of autonomy affects the evaluation of outcomes and so affects the determination of when paternalistic restriction would be for the best.

But the disagreement between constraint-based morality and act consequentialism will often be stark. Suppose that one could prevent two brutal killings of innocent people but only by oneself perpetrating a similar brutal killing of a third innocent person, and there are no other differences between refraining from action in this situation and committing one killing to prevent two identically bad ones. Act consequentialism presumably says one should perpetrate the one killing to bring about overall better consequences; any constraint-based morality that includes a constraint against this type of killing will presumably say one should not. (Suppose the choice is between killing an innocent or not, and the consequences of perpetrating the killing are just ever so slightly better than refraining; consequentialism says one morally must perpetrate the killing.) This implication looks to be counterintuitive. The consequentialist will respond that moral constraints are best interpreted as useful guidelines to help imperfect agents—ill-informed, selfish and biased, and not good at reasoning—make decisions that will tend to be closer to what act consequentialism recommends than those that would be reached by direct application of the act consequentialist principle (Hare 1981, chapter 2).

2g.

Much philosophical writing on what is standardly regarded as applied moral philosophy proceeds on the assumption that moral judgments are likely to be context-specific in some significant way. For example, philosophers investigate bio-medical ethics on the assumption that good health is a unique and special good and that there will be special moral principles that regulate the provision of health care by health care providers and the distribution of health care across the members of a community and the distribution of the broader social determinants of good health across the members of a community or across all human communities. Thus we get books with titles like "Principles of Biomedical Ethics" (Beauchamp and Childers 2008).

This way of thinking shapes Michael Walzer's thinking about social justice in his book *Spheres of Justice*. His idea very roughly is that there are distinct and separate goods in human life. The meaning of a good is set by the beliefs about its meaning by the members of a particular community. The meaning of a good in turn determines how it ought to be distributed and what practices and institutions would be appropriate for its provision and distribution. Each good can be associated with a distinct sphere of social life, and justice is done in a community when goods are produced and distributed in accordance with their social meanings.

One might adopt a version of this sort of view without accepting the conventionalist idea that the beliefs of members of a community about a good determine the nature of that good and the nature of its morally correct distribution for that community.

One area of discourse in which the separate spheres idea has been recently contested is the ethics of killing in war. Current conventions of international law and morality do not hold soldiers who fight for their country in a declared war to be guilty of criminal wrongdoing if their cause is unjust. Instead a line is drawn between deliberately attacking combatants and noncombatants (or bystanders). A soldier fighting in a declared war is (claimed to be) morally permitted to aim at enemy combatants but not deliberately to attack noncombatants even to gain military advantage (Walzer 2015).

This way of thinking has been challenged on the ground that it does not mesh coherently with our broader ideas about the conditions under which it is morally permissible to attack other persons, in self-defense, for example. Jeff McMahan (2004) has argued that one becomes morally liable to attack only if one is responsible for sufficiently wrongful harm to others, the prevention or mitigation of which, to a sufficient degree, attacking one would accomplish. Except in extreme situations one is morally permitted to attack only those morally liable to attack. But if one is fighting in war for an unjust cause, one's acts of violence cannot satisfy this condition, so one's morally permitted options are to flee or surrender. On this view, the ethics of war must cohere with our broader thinking about permissible violence. An ethics of war that is coherent in this way will deny that it is acceptable for enemy soldiers fighting for an unjust cause to kill enemy combatants fighting for a just cause.

When is it morally acceptable to kill another in order to save one's own life or save oneself from grievous injury? Are the moral principles governing the morally permissible deployment of state-sponsored violence different from those that govern the use of violence in other contexts? Is warfare morally special, so that the ethics of war will consist of special principles tailor-made for that sphere of activity? The answers to these questions are disputed, and so far unresolved. As we have seen, this situation is business as usual for applied ethics controversies. If there is progress, it lies more in understanding the structure of disputes and the grounds adduced by proponents of opposed views rather than in resolving the disagreements.

References.

- Beauchamp, T., and Childers, J. 2008. *Principles of Biomedical Ethics*.
- Boonin, D. 2003. *A Defense of Abortion*. Cambridge: Cambridge University Press.
- Feinberg, J. *The Moral Limits of the Criminal Law*, vol. 1 *Harm to Others* (1984), vol. 2 *Offense to Others* (1985), vol. 3 *Harm to Self* (1986), vol. 4 *Harmless Wrongdoing* (1990). Oxford: Oxford University Press.
- Hare, R.M. 1981. *Moral Thinking: Its Levels, Methods, and Point*. Oxford: Oxford University Press.
- Kamm, F. 1992. *Creation and Abortion: A Study in Moral and Legal Philosophy*. Oxford: Oxford University Press.
- Marquis, D. 1989. "Why Abortion is Immoral". *Journal of Philosophy* 86: 183-202.
- McMahan, J. 2004. "The Ethics of Killing in War". *Ethics* 14(4): 693-733.
- Rawls, J., 1999. *A Theory of Justice*, rev. ed. Cambridge, MA: Harvard University Press.
- Ross, W. 1930. *The Right and the Good*. Oxford: Oxford University Press.
- Scanlon, T. 2003. "Rawls on Justification," 139-167, in S. Freeman, ed., *The Cambridge Companion to Rawls*. Cambridge: Cambridge University Press.
- Singer, P. 1972. "Famine, Affluence, and Morality". *Philosophy and Public Affairs* 1(3): 229-243.
- Sinnott-Armstrong, W. 2011. "Consequentialism". *Stanford Encyclopedia of Philosophy*, www.plato.stanford.edu/entries/consequentialism.
- Thomson, J. 1971. "A Defense of Abortion". *Philosophy and Public Affairs* 1: 47-66.
- Walzer, M. 2015. Originally published 1977. *Just and Unjust Wars: A Moral Argument with Historical Illustrations*. New York: Basic Books.
- Walzer, M. 1983. *Spheres of Justice*. New York: Basic Books.
- Warren, M. 1973. "On The Moral and Legal Status of Abortion." *The Monist* 57(1): 43-61.
- Wilson, M. 2006. *Wayward Significance: An Essay on Conceptual Behavior*. Oxford: Oxford University Press.