

Introduction

But is it not monstrous to suppose that if we *have* worked out the consequences and if we have perfect faith in the impartiality of our calculations, and if we *know* that in this instance to break [rule] *R* will have better results than to keep it, we should nevertheless obey the rule? Is it not to erect *R* into a sort of idol if we keep it when breaking it will prevent, say, some avoidable misery? Is not this a form of superstitious rule-worship (easily explicable psychologically) and not the rational thought of a philosopher?

J. J. C. Smart (1956)¹

outsmart, v. To embrace the conclusion of one's opponent's *reductio ad absurdum* argument. "They thought they had me, but I outsmarted them. I agreed that it *was* sometimes just to hang an innocent man." [Satirical reference to utilitarian philosopher J. J. C. Smart.]

The Philosophical Lexicon by Dennett and Steglich-Petersen (2008)

The Metaphor of Legislation in Ethics and Politics

In ethics, one of the most common exercises is to ask, "What would happen if everyone acted that way?" and to then consider what consequences would follow if everyone actually did. There is a potential tension between the two steps. "If" emphasizes the hypothetical, and indeed counterfactual, nature of the thought experiment. My action will not magically cause everyone in the world to act similarly in similar situations. "What would happen," on the other hand, derives much of its significance because we care about the actual

¹ This book covers a four-century timespan, and norms regarding spelling, grammar, formatting, punctuation, and capitalization have changed quite a bit over the years. I have chosen to reproduce the original texts in quotations as they appear in my sources except that I have silently removed some extra spaces. All italics are from the original source unless otherwise noted. Where these differ from modern usage, they are a reminder that we are in conversation with voices from a different era but with whom we can still converse.

consequences that will affect the lives of actual people, not hypothetical or counterfactual consequences. If we think that our answer to this question helps determine what is right for ourselves and for others, our question can be reframed in legislative terms: “If you were legislating a code of conduct for all to follow, what code would be best?” In this book I will explore the historical reasons why this common and intuitive way of thinking about ethics came to be regarded as morally problematic and I will suggest a less paradoxical way it could be incorporated into contemporary moral deliberation in ethics and in politics.

We can say roughly that the legislative perspective is one where deliberation is aimed at rule selection: “What general rule would be best for situations like this one?” In stating it this way, I am consciously excluding some other sorts of activities and considerations that are relevant to real legislators. In a democracy, a legislator might weigh the impact on their reelection chances of voting yes or no on a piece of legislation in terms of reactions of voters, campaign contributors, and the media. They might need to consider trading their vote on the proposed legislation to gain votes for a different bill that they think is more important. I am excluding these sorts of considerations insofar as they are simply prudential questions about what is best for the legislator but I include them insofar as they relate to the public good. A legislator who believes law A would be best, but who also believes that there is little chance of A being adopted, might choose to propose law B instead, which would still be a substantial improvement over the status quo, though not the best, especially if proposing A would decrease the chances of adopting B. This sort of thinking also counts as use of the legislative perspective since it is still oriented toward an action-guiding rule. It aims at the best attainable rule.

My particular interest is in the use of this deliberative frame outside of the obvious context of actual legislators deliberating on the merits of actual laws. It can be used in quasi-legislative activities, or as a heuristic for helping resolve ethical dilemmas of individuals, or as a source of moral guidance and constraint for those who have more traditional legislative power. The following contemporary example will illustrate what I mean by the legislative point of view and the moral questions that arise when we consider using it in some contexts as opposed to others.

Ethics and the Legislator: Stepping Out of the Chamber

Imagine a thoughtful legislator sitting in the legislative chambers thinking about whether to support a proposed piece of legislation. The legislation, if passed, would give more latitude to immigration officials to grant asylum

Introduction

3

requests. In deciding whether to support the bill, they will think about the probable consequences of its adoption. If more asylum seekers are accommodated, how much better (and longer) will their lives be compared with returning them to their home countries? If the law increases the number of asylum seekers without increasing the funding for investigating and processing their cases, will the result be longer waiting times for applicants? What implications will there be for the economy of the legislator's own country, including the welfare of its workers? If discretion increases for immigration officials, what if they use that discretion to the benefit of applicants from some countries rather than others? What if the use of discretion is affected by racial bias? What if the latitude leads to more mistakes and, as a result, dangerous criminals are accidentally let into the country? The legislator is particularly moved by the plight of asylum seekers from a specific country in the midst of a religiously motivated civil war and wishes the law made it easier for immigration officials to grant asylum in such cases, but they also know that the law would apply to applicants from all countries, not just that one.

All of these questions raise ethical issues and there are a variety of ethical frameworks the legislator might use to decide what they ought to do with regard to their vote on the proposed legislation. One approach is consequentialism, the view that consequences alone determine what is right and wrong. The legislator has a sense that some outcomes are better than others and that in at least some cases they can weigh these against each other so that they can determine whether the new legislation or the status quo is more likely to lead to better consequences. They might look at the actual frequency of crimes committed by those granted asylum compared with the rest of the population and compare that with the likely harms experienced by those denied asylum, and decide that the consequences of the proposed bill would be, on balance, an improvement. What is of particular interest is that the sort of decision the legislator is faced with (to support proposed legislation or not) gives them a particular perspective from which they think about those consequences. They must think of the consequences of a new law being adopted, not just of a particular set of worthy applicants being admitted. The new law will give discretion to people other than the legislator and they have to think about how these others will likely use the discretion rather than about how they would use the discretion if they were making the decisions about specific applicants. Those who will interpret and apply the law will be fallible and make mistakes. The new law, once known, may influence the behavior of those potentially affected by it.

Now suppose that our legislator steps out of the chamber after voting in favor of the bill. They receive a call from an important campaign contributor who wants the legislator to use their influence to help them secure a lucrative government contract. Suppose the legislator thinks consequentialism informs them regarding which action is right in this case as well. They might reason that a general practice of legislators helping donors in this way would have overall negative consequences. The government could end up with inferior value for the contract and public trust in the fairness of the system would be undermined. There would be obvious beneficial consequences for the legislator in keeping the contributor happy. Let us suppose the contributor would likely provide reasonable value to the government in return for the contract. These situation-specific benefits are vastly outweighed by the negative consequences of a political system where legislators systematically use their influence to get governmental contracts for campaign contributors or their friends and family. The legislator might, therefore, have good consequentialist reasons to say no.

Or they might not. They might instead, as a consequentialist, assess the chances that use of their influence will become publicly known and conclude that the chances are very small. They might, therefore, assume that the chances that their action will contribute to the formation of a new norm that then influences other legislators or public perceptions is also quite small. Looking only at the consequences of this one particular decision, they might think about the benefits to their reelection chances that would come with keeping the contributor happy and the important legislative causes (like helping refugees) that they think they will be able to advance if they continue to serve. If there is little reason to believe that there will be a significant causal link between their actions and those of other legislators or public perception, they might question why counterfactual consequences matter. What matters instead are the probable consequences of this specific action.

In the two scenarios above we see a legislator confronting ethical dilemmas and trying to use consequentialism to resolve them. In the first scenario, the context in which the legislator must make their decision (voting on potential legislation) structures their consequentialist thinking in a particular way that focuses on the consequences of the adoption of a publicly known rule that will be interpreted and enforced by fallible, biased people. From their view in the chamber, questions about the probable consequences of a publicly known rule are not counterfactual since the result of passing the legislation is the adoption of a new public rule.

Introduction

5

In the second scenario the legislator leaves the chamber. Although they could take a situated perspective and look only at the morally relevant features of the specific acts they might perform, many people find it ethically appealing to use something like the legislative perspective outside the chamber. Many people's moral intuitions resonate with something like the golden rule with its demand that our actions be consistent with reciprocity. We should not do what we would not want others to do, and imagining a general public rule that would approve of our actions is a way to think about our actions as reciprocal. Moreover, even if many people think that consequences are not the whole of ethics, few people think them irrelevant to ethical decision-making. John Rawls remarked that "All ethical doctrines worth our attention take consequences into account in judging rightness. One which did not would simply be irrational, crazy" (Rawls 1971, 30). It is possible to account for many widely held moral and ethical beliefs by noting that a given rule regulating our actions would produce better consequences than alternative rules or no rule at all. If a rule would be beneficial if adopted and followed by others, that gives me a moral reason to act on it.

This line of thought, though intuitively attractive, has been found deficient, at least insofar as it is thoroughly grounded in the most well-known form of consequentialism, utilitarianism. Although the definition of utilitarianism is disputed, we can for now define it as an approach to ethics that judges the rightness of actions by the likely or actual impact of those actions on happiness (or pleasure, welfare, or satisfied preferences). Consequentialism is a broader term referring to approaches that define actions as right or wrong based on whether those actions produce good outcomes but that may have a broader or more complex account of what counts as a good outcome. A lot of ink was spilled in the twentieth century debating the merits of "act-" and "rule-" utilitarianism, and this debate tracks along with the two scenarios above. Rule-utilitarians argued that one should not apply the principle of utility ("maximize happiness") directly to specific, situated actions. Instead, one should use the principle of utility to select rules for behavior and decide specific actions with respect to those rules. This indirect form of utilitarianism can more easily explain why the legislator should not use their influence to help a donor get a government contract. There was a small cottage industry generating hypothetical examples where act-utilitarianism seems to lead to troubling ethical conclusions, especially if one believes one's actions can be kept secret (framing an innocent person to prevent a riot, harvesting organs from a healthy patient without their consent to save the lives of five other patients, etc.). If

one looks only at the specific action, one can construct a scenario where utilitarianism seems to compel unethical action. We can add to this cases where the harms or benefits of a single act are negligible but where a general practice of many people performing that act could lead to real harms (not voting, walking on the grass of a pristine lawn, etc).²

This rule-utilitarian line of thought was found wanting because it seemed inconsistent with its utilitarian foundations. If I care about maximizing happiness it seems I should care about the probable consequences of *this* action and not worry about counterfactual consequences that would occur if I were legislating a rule for everybody, given that I am not actually legislating a rule for everyone (Lyons 1965). Put another way, a utilitarian has no trouble explaining why an actual legislator should utilize a legislative perspective when actually legislating. If you really are enacting a public rule, think about the impact on happiness of adopting that rule. What is puzzling is why a theory committed to producing the best actual consequences would tell you to think like a legislator when you are not legislating.

The Historical Emergence of Counterfactual Use of the Legislative Perspective

This puzzle was at the center of the twentieth-century debate between “act-” and “rule-” utilitarians, a debate nurtured in part by some interpreting Mill as a rule-utilitarian (Urmson 1953). The quote from Smart (1956) in the epigraph summarizes in three devastating questions the basic problem: to obey the rule in cases where you could produce better consequences by breaking it seems to treat the rule idolatrously, even worshipfully, rather than rationally and philosophically, since the rule has no inherent authority of its own, it is only a pointer to what normally produces good consequences. The second epigraph shows the problem for act-utilitarians like Smart: he really is forced to say that the rule “don’t use your discretion to bring about the execution of a person you know to be innocent” should be broken in exceptional cases when doing so produces more utility than following the rule. It is such counterintuitive conclusions that have often motivated the

² On framing the innocent see Rawls 1955. Early discussions of doctors killing patients to harvest their organs are found in Harris 1975 and Harman 1977. Though not originally conceived of as a contribution to the debate about act- and rule-utilitarianism, nonconsensual organ harvesting became a typical instance of something that is not utility maximizing as a general practice though it might be in particular instances. Harris’s survival lottery proposal arguably could be affirmed by rule-utilitarians, not just act-utilitarians. Harman’s version (page 3) is the one that became canonical.

desire to treat rules with more reverence. In this book, I will argue that Smart's linking of rule-utilitarianism to religion is telling, probably more telling than he knew. The instability of rule-utilitarianism can be better understood if situated within a different historical frame that shows how the secular, realistic, consequentialist, and counterfactual rule-utilitarians of the twentieth century are the heirs of a theistic, realistic, weakly consequentialist stream of thought that was not thought, by its practitioners, to be counterfactual at all.

What I have called the legislative perspective descends from a particular way of thinking about what is right and about justice that was popularized by thinkers whose starting point was a divine legislator motivated by benevolence. In its original formulation, people might think about what rules would be rational from God's perspective in order to discover the laws they were to follow, but they would not have thought of themselves as the legislators of the moral law. While they used a legislative lens to reason to their conclusions, the heuristic was a way of understanding what was independently true for others as well, rather than a decision-making strategy for themselves alone. For them, the move from the situated perspective to the legislative perspective was not hypothetical or counterfactual as they assumed the existence of an actual God promulgating actual moral laws. They also assumed a moral obligation to obey God's law. Their focus on morality in terms of obligation to obey law is characteristic of much of modern moral philosophy (Anscombe 1958).

Since I am ultimately interested in hybrid approaches, I will use the phrase "weak consequentialism" to designate approaches that adjust the content of rules and principles to produce better consequences but that can also include nonconsequentialist commitments. The term "consequentialist" was not in use when these authors were writing so they were not participating in a defined movement of that sort. The theological character of many of their works would keep them from being classified as thoroughgoing consequentialists today by many people. They might be better classified as hybrid theorists in that their moral thinking relied upon nonconsequentialist theological commitments to frame their consequentialist reasoning. I will use the term "hybrid" in this book as a broad category for theories that include both consequentialist and nonconsequentialist elements. Many of these thinkers, for example, held to a belief that God's moral law should be obeyed, which was not explicable in purely consequentialist terms while using consequentialist reasoning to fill out the content of the moral law. Part II of this book will argue for the viability of a modern hybrid approach that is available to people who reject the theological assumptions of the earlier approaches.

In the seventeenth and eighteenth centuries this theological tradition of morality as the will of a divine lawgiver was increasingly separated from its original biblical context and joined to the idea of a benevolent divine lawgiver. Important figures in this line of thought include Hugo Grotius, Richard Cumberland, John Locke, Francis Hutcheson, William Paley, and John Austin. Their more religiously-minded critics at the time complained that God's holiness or justice had been lost due to an overemphasis on God's benevolence. The excitement over scientific theories that could explain natural phenomena by means of "laws of nature" may have encouraged these authors to try to find a similar approach in ethics. If we start with the assumption that there is a moral code willed by God that defines what is right and add the assumption that the overriding goal of that God is bringing about good outcomes, one has a theological version of the rule-consequentialist position. The moral rules need not be discovered from the Bible. Instead, one can determine the correct moral action by asking what moral code a benevolent God would promulgate to human beings. One can even stipulate that God, as a benevolent legislator, must attend to the same kinds of considerations as a human legislator. The divine moral code's content must account for the selfishness, fallibility, and other limitations of the mere mortals on whom it is imposed. In other words, the figure of the divine legislator allowed theorists to develop a weak version of rule-consequentialism that, unlike its secular counterpart, was not, from their perspective, counterfactual. In the religious version it is assumed that there is in fact a benevolent legislator who has enacted a welfare-maximizing moral code.

In the nineteenth century, John Stuart Mill and Henry Sidgwick, building on the earlier thought of David Hume and Jeremy Bentham, worked to figure out what the standing of rules could be in a thoroughly secular and utilitarian philosophical system. Their solutions came with costs that many subsequent philosophers were unwilling to pay. In the twentieth century, there was a resurgence of explicitly secular, rule-utilitarian thought that tried to use legislative reasoning to avoid some of the counterintuitive conclusions of act-utilitarianism. These philosophers were often criticized, along the lines of Smart's quote in the epigraph, for being inconsistent utilitarians. In fact, a number of the prominent twentieth-century philosophers who are often thought of as rule-utilitarians are better characterized as adopting hybrid approaches that are, in a way, descendants of the hybrid theological approaches of earlier centuries. Hybrid approaches represent the best opportunity for continued use of the legislative perspective in cases where the agent is not literally legislating.

Different Forms of the Legislative Perspective

So far, I have been talking about the legislative perspective as if it is just one way of looking at things when, in fact, it is a family of ways. (I use the phrases “legislative perspective” and “legislative point of view” interchangeably.) In this section, I describe how conceptions of the legislative perspective can vary across four dimensions. In the next section I will describe four contexts in which the legislative perspective could be employed. The historical study will reveal differences, not just similarities. My interest in Part II is in asking how a particular version of the legislative perspective is justified. It is thus important both for framing the historical inquiry and for thinking about the contemporary relevance of the legislative point of view to specify the different forms the legislative perspective can take. Our present concern is with the sort of deliberation one engages in when thinking legislatively, that is, the sorts of considerations that are morally relevant. This is separate from the question of the contexts in which it is appropriate for a person to employ the legislative point of view, however construed. That will be addressed in the following section. What all of the variations of the legislative perspective discussed here have in common is that they are about the selection of action-guiding rules that apply to a whole set of cases. By contrast, a situated perspective (the main alternative to the legislative perspective) is interested not in identifying the rule that would be best to direct a set of cases but in identifying the action that would be best in this particular situation.

The first dimension along which the legislative perspective can vary is how realistic or idealized legislative deliberation is. By “realistic” I mean how closely does it track with the considerations a legislator would typically use in deciding whether to vote in favor of proposed legislation. Immanuel Kant wrote, “*Act only in accordance with that maxim through which you can at the same time will that it become a universal law*” (Kant G, 31). In another formulation he wrote, “act in accordance with the maxims of a member giving universal laws for a merely possible kingdom of ends” (Kant G, 46). While Kant was open to thinking of the moral law as, in a sense, the will of a supreme being, duty to God “is not objective, an obligation to perform certain services for another, but only subjective, for the sake of strengthening the moral incentive in our own lawgiving reason” (Kant M, 230).³ Kant’s formulation, as commonly interpreted, is explicitly

³ Kant also wrote, “A law that binds us *a priori* and unconditionally by our own reason can also be expressed as proceeding from the will of a supreme lawgiver, that is, one who has only rights and no duties (hence from the divine will); but this signifies only the idea of a moral being whose will is a law for everyone, without his being thought of as the author of the law” (Kant M, 19).

legislative and highly idealized. I need not worry about whether my opinion is the one that will prevail in the legislative chamber since I can act as if my maxim once adopted is a universal law. I legislate for an idealized kingdom of ends where people act in accord with their duty even though I know such a kingdom is “merely possible” and not actual.

By contrast, consequentialist approaches that tend toward a more realistic conception of the legislative point of view would consider it relevant that the best option on the merits might be one that few people would adopt and that this might be a relevant reason for rejecting it. It is similar to the way that a legislator might decide that it is better to support a law that has a higher probability of becoming law even though it is not the best law that could conceivably be adopted. A realistic approach would consider how fallible and sometimes self-interested people would interpret and apply the law and might reject a proposed law because of foreseen errors. Legislators, using the realistic approach, are not the ones who interpret and apply the law and must anticipate how other people with different values and judgment will interact with the law. They must also consider the likelihood that people can be brought to comply with a law and the costs involved in attempting to bring about compliance, including voluntary compliance. My interest in this book is primarily in the more realistic theories that more closely mirror the deliberation of human legislators. Part II will explore ways in which Kantian approaches can be combined with this sort of realism.

A second dimension along which use of the legislative perspective can differ is with respect to consequentialism. How does a legislator judge which law would be best? A pure consequentialist would take the expected outcomes of the various options to be the only relevant consideration. Kant insisted that legislators should not worry about the likely consequences of their laws when legislating (Kant *M*, 109). John Stuart Mill claimed that even Kant was a consequentialist in the end, since Kant rejects maxims by considering the consequences of the universal law (Mill *CW*, 1:207). Most have thought Mill was wrong about Kant, but Mill was certainly right about real legislators. One can hardly imagine them being indifferent to the likely consequences of a proposed law, especially since “consequences” is here used very broadly and could include things like the outcome of more people’s rights being protected. In this book, my primary interest is in theories that take the likely consequences of different frameworks of rules to be morally relevant to deciding whether those rules rightly influence our decision-making and our understanding of what is right and wrong. Theories need not be wholly consequentialist, but they must be at