

RESEARCH ARTICLE

# Leveraging a Sturdy Norm: How Ethicists Really Argue<sup>1\*</sup>

David DeGrazia

Department of Philosophy, George Washington University, Rome Hall 566, Washington, DC 20052, USA  
Email: [ddd@gwu.edu](mailto:ddd@gwu.edu)

## Abstract

Rarely do everyday discussions of ethical issues invoke ethical theories. Even ethicists deploy ethical theories less frequently than one might expect. In my experience, the most powerful ethical arguments rarely appeal to an ethical theory. How is this possible? I contend that ethical argumentation can proceed successfully without invoking any ethical theory because the structure of good ethical argumentation involves leveraging a sturdy norm, where the norm is usually far more specific than a complete ethical theory. To illustrate this idea, I present the argumentative structure of five powerful articles in the ethics literature. I further argue that the present model of ethical argumentation is consistent with the coherence model of ethical justification, but the former need not—and usually should not—invoke the latter explicitly for various practical reasons.

**Keywords:** ethical theory; reflective equilibrium; coherence; leveraging; norm; application; considered judgments

## Introduction

Those of us who have taught ethical theory—in the classroom, in clinical settings, or elsewhere—may have had the experience of wondering what the relationship between ethical theory and ethical argument *really* is. The question is motivated by the observation that ethical theories are, in practice, invoked much less often than one might think in light of a very natural understanding of the role of these theories: to be *applied* in order to generate justified answers to particular ethical questions. Rarely everyday discussions of contested ethical issues invoke ethical theories. Moreover, even ethicists deploy ethical theories less frequently than one might expect. In my experience, the most powerful ethical arguments *rarely* appeal to an ethical theory. How is this possible?

Sophisticated ethicists might suggest that the reason ethical theories are infrequently deployed, or needed, in practice is that the model of applying these theories is inferior to the subtler coherence model of ethical justification (also known as reflective equilibrium). Yet, even the coherence model is seldom mentioned or explicitly invoked in ethical argumentation. How can this be if it is really the best model of ethical justification?

I contend that ethical argumentation can proceed successfully without invoking any ethical theory because the structure of a good ethical argument involves *leveraging a sturdy norm*, where the norm is usually far more specific than a complete ethical theory. I further argue that this process is consistent with the coherence model of ethical justification but need not—and usually should not—invoke it explicitly for various practical reasons.

In the sections that follow, this paper (1) presents the application model and explains why it rarely captures the structure of ethical argument; (2) identifies several argumentative strategies represented by articles in the ethics literature that have been lauded for their argumentative power; (3) demonstrates that each of these argumentative strategies involves leveraging a sturdy norm; (4) clarifies the relationships

---

\*The online version of this article has been updated since original publication. A notice detailing the change has also been published.

among my model of ethical argument, ethical theories, and the coherence model of ethical justification; and (5) closes with concluding reflections.

What will this discussion add to the literature? Readers well-versed in ethical theory, “applied” or “practical” ethics, and methodology in ethics might find much of what I say about the application and coherence models to be familiar. But others—including health professionals, philosophers working outside of ethics, and students—might find these discussions genuinely informative. As for my model of ethical argument, I expect that many others have had similar ideas—yet I have never seen this model presented explicitly, much less with the metaphor I use to make it vivid. I am also fairly sure that some of the connections I draw between the coherence model of ethical justification and my model of ethical argument are novel.

### The Application Model: Applying an Ethical Theory

In the application model, the appropriate way to address a particular issue in normative ethics is to apply an ethical theory to the problem area to yield a verdict. An ethical theory is a structure of moral norms that is intended to offer guidance in particular instances of ethical decision-making, to be as general as possible, and to be more or less comprehensive of the ethical domain.<sup>2</sup> Naturally, since there are multiple ethical theories on offer, according to the application model, one should apply the *correct or most adequate* theory in order to have justified confidence in the result. An example will serve to illustrate both the use of this model and its considerable limits.

In teaching ethical theory, faculty often begins with utilitarianism, both because of its historical prominence in anglophone ethics and because its most classic version is relatively easy for students to understand. So let us consider direct hedonistic utilitarianism, which features this principle of utility: the right action is that which, among those available to the agent, has the greatest expected balance of happiness over unhappiness, where everyone’s interests have equal weight.<sup>3</sup> Put simply, this theory directs the agent to “maximize happiness.”

An ethical theory should help us resolve ethical quandaries such as whether we should permit euthanasia in medical settings. Pioneering bioethicist Joseph Fletcher deployed the above principle of utility in arguing that, indeed, we should permit euthanasia: much needless suffering would thereby be averted, scarce medical resources would be better conserved, and no rational moral rule would be violated.<sup>4</sup> Fletcher’s reasoning provides a nice example of applying an ethical theory to a controversial issue in practical ethics and generating a significant result: a liberal approach to euthanasia.

As this example suggests, ethicists do sometimes apply ethical theories to address particular ethical issues. But they do so less frequently than one might expect. Why? The answer concerns several important limitations of the application model.<sup>5</sup>

Consider again direct hedonistic utilitarianism. Even within the utilitarian camp, it is controversial: many utilitarian theorists today reject the hedonistic value theory (utility or well-being construed as happiness)<sup>6</sup>, the direct nature of the action guide (instructing agents to deploy the principle of utility directly rather than comply with rules that generally promote utility)<sup>7</sup>, or both. Furthermore, most ethicists are not any sort of utilitarian. Some are consequentialists who, like utilitarians, understand right action in terms of expected consequences but disagree with utilitarians that right action is simply a matter of maximizing utility.<sup>8</sup> Many other ethical theorists are deontologists, who like consequentialists, have a criterion—or a set of criteria—for right action but, unlike consequentialists, deny that this standard can be cashed out entirely, if at all, in terms of expected consequences.<sup>9</sup> (Deontology covers a massive range of theories, ranging from Kantian ethics to Rossian ethics, from Rawls’ theory of justice to classic libertarianism.) Then, there are virtue ethicists, whose approach to ethics does not focus on right action at all but rather on moral character, though some virtue ethicists hold that criteria of right action may be derived from considerations of character.<sup>10</sup> There are other ethical theories as well, such as the ethics of care<sup>11</sup> and the capability approach<sup>12</sup>; whether these are genuine alternatives to consequentialism, deontology, and virtue ethics or instead can be located within this penumbra, is unimportant here. What is important is that there are quite a few ethical theories on offer.

Given the plurality of ethical theories, no one of which commands majority assent—much less consensus—among experts or the lay public, which ethical theory should one deploy? If all respected ethical theories generated the same implications on contested ethical issues, then while the theories themselves might be unnecessary, one could simply apply any one of them as a reliable method of getting a justified answer. But different ethical theories often have different ethical implications, so the choice of theory matters. Classic libertarians and most consequentialists will disagree on optimal gun policy, for example, with the former vindicating permissive policies and the latter favoring much more substantial regulation.

An open-minded ethical thinker who is committed to the application model should be troubled by the plurality of respected ethical theories. It is unclear which ethical theory one should apply. Now, if it seems clear *to her* which theory is best—she has expertise in ethical theory, has carefully considered the best ones, and has judged Theory X to be most adequate—she might feel confident that the way to do ethics well is to apply Theory X wherever there is moral uncertainty. Yet, two problems remain even for this theoretically confident moral agent.

First, ethical theories tend to be highly general and abstract such that their implications in particular contexts are often uncertain. That is, an ethical theory can be indeterminate. Take, for example, Kantian ethics, whose supreme principle, the Categorical Imperative, directs us to act in such a way that one always treats human beings as ends-in-themselves (beings with dignity) and never uses a human being merely as means to anyone's ends (which would be treating that person as a tool). This principle captures a powerful ethical idea with which even non-Kantians tend to resonate to some extent. But now try to apply this idea to an issue that confronts Emergency Departments: Should priority be given to those patients who arrive earlier, those who are in greater medical need, and those who are more likely to benefit—or should other criteria be used? Treating patients as ends-in-themselves and not using them as mere means is too abstract a standard to inform us here. So, even a moral agent who has great confidence in an ethical theory will sometimes find that attempting to apply it is not the best way to answer a specific ethical question.

A further difficulty confronts the confident ethical theorist: she must sometimes try to persuade other people about what action or policy would be right, and rarely will all of her interlocutors or audience share her confidence in Theory X.<sup>13</sup> How can she convince others—in informal discussion, in clinical rounds or Grand Rounds, or in a journal article or book devoted to some problem in practical ethics? She could try to “prove” Theory X, but such an effort would be time-consuming and fairly unlikely to succeed.

In sum, the application model is severely limited as a method for justifiably resolving concrete moral problems. First, it is far from clear which ethical theory is the most adequate. Second, any given ethical theory is probably indeterminate in some contexts. Third, many of one's interlocutors might not share one's commitment to a particular ethical theory, making persuasion by application of the theory unlikely.

No wonder ethicists do not apply ethical theories very often, especially when they are attempting to persuade others on a particular issue in normative ethics.

### Several Argumentative Strategies

Let us consider several articles that illustrate argumentative strategies, or methods, other than applying an ethical theory. Each of the articles has been widely cited and most have been reprinted in anthologies. Let me add my impression that each is quite powerfully argued (whether or not it ultimately persuades me of its conclusion).

Before proceeding to these argumentative strategies, it is worth underscoring that this paper is interested in the structure of *good* ethical arguments rather than the structure of *all* ethical arguments. Many ethical arguments are of poor quality. Consider, for example, an argument that begins with the assertion that, *from a Marxist perspective*, hospitals should be structured in such-and-such a way and then goes on to suggest that, for this reason, hospitals should be structured in this way. This is a bad argument because, in effect, it appeals to the supposed authority of Karl Marx, fails to recognize the considerable reasons one might disagree with Marx's views, and ignores the fact that, rightly or wrongly, most people are not Marxists. Note that one could produce an equally bad argument by beginning with

“from a classic libertarian perspective . . .,” implicitly appealing to the authority of Adam Smith or Milton Friedman, and ignoring the considerable reasons one might have to reject classic libertarianism.<sup>14</sup>

Instead of attempting to display the structure of all ethical arguments, I will focus on the structure of arguments that do not commit obvious logical fallacies or other errors and sometimes succeed in changing the minds of thoughtful people. These are good arguments in the sense that they are both high-quality and capable of persuading people. They are the sorts of arguments ethicists should want to produce.

Argument by analogy. In one of the most highly cited and anthologized philosophy articles of all time, Judith Jarvis Thomson argues for the ethical permissibility of abortion in many or most cases.<sup>15</sup> Her chief arguments are two arguments by analogy. In the more famous of the two, you are asked to imagine being captured and involuntarily placed in a hospital, where you are expected to remain for nine months for the sole purpose of keeping a famous violinist alive through the use of your kidneys. In the less famous but equally important argument, you open your window for fresh air and, as you know could happen (and despite your using a screen), a “people-seed” enters, implants in a piece of furniture, and begins to grow. Both thought experiments are designed to elicit the response that it would be permissible for you to part ways with the innocent person who needs you—by disconnecting from the violinist or evicting the dependent people-seed from your house. Thomson argues that, in many or most cases, aborting an unwanted fetus is analogous in ethically relevant respects.

Appeal to an intuitively plausible norm. In another of the most highly cited and anthologized philosophy articles of all time, Peter Singer suggests a plausible moral principle: if we can prevent something bad from happening without sacrificing anything of moral importance, we should do so.<sup>16</sup> He then argues at length that most of us can contribute considerable funds to help those at risk of famine without sacrificing anything very important, supporting his thesis that we have substantial duties of beneficence to help those in greatest need. This conclusion, he notes, is at odds with conventional morality, which treats charity as optional rather than obligatory.

Appeal to considered moral judgments in support of a theory, which is then applied. Tom Regan’s book on animal rights is a classic contribution to animal ethics.<sup>17</sup> Condensing much of the book’s content, the present article succinctly invokes considered moral judgments to support two claims: that animals have inherent value (or moral status) and that utilitarianism is an inadequate ethical theory.<sup>18</sup> Regan uses these two claims to motivate his animal rights theory, according to which “subjects of a life” have rights that are not to be overridden even in the name of utility. He then applies this theory to support abolitionist positions regarding animal husbandry, animal research, and animal trapping. What strikes me as impressive about this article is not the (rather straightforward) application of the animal rights theory but rather the skillful appeal to considered judgments to motivate the theory.<sup>19</sup>

Conceptual analysis to elucidate the content of accepted norms. In an article that was excerpted for at least one anthology, Ruth Macklin notes the common perception that respect for patient autonomy and respect for cultural differences are sometimes incompatible in medical settings because some patients represent cultures that do not prize individual autonomy as highly as Western society does.<sup>20</sup> Macklin carefully examines respect for autonomy and shows that it does not, as some commentators suggest, support (1) *requiring* people from traditional cultures to make medical decisions for themselves or (2) *imposing* unwanted medical test results on them. She also argues that *reasonable* respect for cultural differences does not always defer to a culture’s norms (for example, norms favoring female genital mutilation). She concludes that respect for autonomy and respect for cultural traditions, properly understood, are compatible.

Factual illumination in light of shared moral goals. Emphasizing the moral importance of two widely shared goals—universal access to health care and limiting healthcare expenditures—Steffie Woolhandler and David Himmelstein cite empirical data about the American and Canadian healthcare systems to argue that, contrary to popular American belief, the United States could achieve universal access without increasing its healthcare budget.<sup>21</sup> So impressive was their argument that, despite its leftist political orientation, it was published by the relatively conservative medical journal *JAMA* long before Democratic politicians were publicly supporting what is now called “Medicare-for-all” or a “single-payer system.”

The presentation of argumentative strategies in this section is not meant to be exhaustive. Rather, it illustrates several types of argumentation deployed in high-quality ethical discussions that have proved capable of changing some readers' minds without assuming any ethical theory at the outset. What do these argumentative strategies have in common?

### Leveraging a Sturdy Norm: The Structure of (Good) Ethical Argumentation

The five argumentative strategies illustrated above have a common general structure. In each case, one uses a "sturdy" or confidently held norm—an ethical principle, rule, moral goal, or one or more considered judgments—to reach a significant result, one that might be surprising or is less commonly held than the sturdy norm used to reach it. The norm is leveraged by applying a sort of argumentative pressure, supplied by novel ideas, reasoning, or information, explaining how it is possible to move from a confidently held norm to a surprising one.

A visual image might help here. Picture a seesaw at a children's park. A small child sits on one side of the seesaw, which is touching the ground. Think of the child as a currently stationary ethical belief. Think of the seesaw as a study norm, which can be shared by both the child and someone else. Imagine that a heavier child comes over and either pushes down on the side of the seesaw that is elevated or sits on it. The big child's downward pressure on the sturdy seesaw is like argumentative power: new ideas, novel reasoning, or previously unknown empirical information. Using this novel force, the shared sturdy norm is leveraged to induce a change of position on an ethical issue (just as the spatial position of the smaller child changes).

Let us now consider each of the arguments that served as examples.

In Thomson's article, the sturdy norm in the more famous argument is the considered moral judgment that it would be permissible for you to detach from the violinist in the hospital, even though doing so would result in his death. The new idea, which leverages the sturdy norm, is that your situation vis-à-vis the violinist is relevantly similar to a pregnant woman's situation in the case of an unwanted pregnancy. In the less famous argument, the sturdy norm is that it would be permissible for you to remove the growing people-seed from your house, though doing so would entail its death. The new idea is that your situation here is relevantly similar to that of a woman with an unwanted pregnancy. In both arguments, the significant result is that abortion is permissible in a broad range of cases—even if the fetus is a moral person (someone with full moral status).

In Singer's discussion, the sturdy norm is the principle that, if you can prevent something bad from happening without sacrificing anything morally important, you should do so. The novel idea is that most of us—at least those of us reading his article—are in a position to contribute funds that could save lives without sacrificing anything morally important. The significant result is that we are obligated to do so.

In Regan's argument, the sturdy norm is our considered moral judgments about various cases. The novel input is an analysis of these judgments in defense of an animal rights theory, which is itself a significant result and, when applied to various human uses of animals, leads to further significant results about the appropriateness of abolishing such practices.

In Macklin's discussion, there are two sturdy norms: general and somewhat vague principles of respect for autonomy and respect for cultural differences. (Or we could say there is just one sturdy norm: the conjunction of these two principles.) The novel reasoning is the careful analysis and specification of these concepts in light of various cases. The significant result is that the two general norms, which are often thought incompatible in some contexts of medical decision-making, are compatible.

In Woolhandler and Himmelstein's article, the sturdy norms are the goals of universal health care coverage and cost controls. The novel input is empirical evidence about the American and Canadian healthcare systems. The significant result is the compatibility of the two norms and the desirability of the United States's adopting a healthcare system that is financed in a similar way to Canada's.

In none of these five argumentative strategies does one need to assume the correctness of a particular ethical theory—although, in one of them, the correctness of an ethical theory is supported by appeal to considered judgments. I have argued that most good (high-quality and persuasive) ethical arguments do

not assume the truth of an ethical theory. Given the disadvantages of making such an assumption, as discussed earlier, it is usually preferable to identify a more specific norm that is more likely to be accepted by one's interlocutor or reader, and apply leverage to that norm with some new idea, novel reasoning, or fresh information, thereby arriving at a significant ethical result.

Note, however, that even those who are very confident about a particular ethical theory, and who begin an argument by assuming its correctness, are *treating* that ethical theory as a sturdy norm. This is true whether one is attempting to persuade others or is simply making a personal decision about what to do in some ethically significant context. Consider an example of each possibility.

Someone who embraces the tradition of Catholic natural law ethics—a type of deontology—might treat this ethical system as a sturdy norm in discussing whether Nation N's entry into a particular war was just. If this person's interlocutors embrace the same ethical outlook and the just war theory it encompasses, the chances of persuasion are generally greater than if they do not. If the interlocutors do not share the Catholic moral outlook, our theorist might do better to emphasize particular ideas in Catholic just war theory—for example, that nations should not intentionally kill noncombatants—that are fairly likely to be shared by people whose comprehensive moral outlooks are different. In doing so, this person would treat the more specific moral claim as a sturdy norm.

If you are convinced that Theory X is correct, then you might consult it in making personal ethical decisions regarding whether to eat animal products, how much to contribute to charity, or whether to purchase an electric vehicle. To the extent that such decisions are private (and do not involve negotiating with a life partner, say), it does not matter whether other people share your commitment to Theory X. In any case, you would be treating this theory as a sturdy norm, consistent with my thesis that good ethical argumentation involves leveraging a sturdy norm. Here, the argument would be intrapersonal rather than interpersonal.

### How this Model Relates to Ethical Theories and to the Coherence Model of Ethical Justification

An ethical theory offers a structure of general norms for addressing ethical issues. As we have seen, though, warranted confidence in a single ethical theory may be hard to come by. Yet, as have also seen, this uncertainty about theory does not impede ethical argument. Why not? After all, ethical theories are supposed to capture the general structure of *correct* ethical principles or other norms. Does it make sense, then, that good ethical argumentation typically invokes a sturdy norm that is more specific than a theory? How can we be justified in believing that a particular norm is “sturdy” if we do not know which full-blown theory is correct or most adequate?

The answer is that *we have more justified confidence in some ethical judgments than we do in any theory*. That intentionally harming others tends to be wrong is far more obvious than any particular ethical theory. That no legal regimen should permit the enslaving of human beings is similarly more obvious than any ethical theory. That keeping one's promises and speaking truthfully are generally right needs no support from a theory. We may say the same about the judgment that generosity is a virtue. No wonder, then, that plausible ethical theories have considerable overlap in the content of their implications. Indeed, when we identify and leverage a sturdy norm more specific than a theory in advancing a high-quality and effective argument, that norm is likely to lie within the overlap of plausible ethical theories.

That we sometimes have more justified confidence in specific beliefs than in any theory that might account for them is a plausible thesis in at least some domains outside of ethics. For example, I have a highly justified belief that, as I write these words, I have a laptop on my lap. By what theory can I justify this perception-based belief? Is direct realism true: do I directly perceive the material (mind-independent) objects that are my laptop, hands, and lap? Is indirect realism correct: do I perceive these material objects indirectly, and directly perceive only certain mental events? Or is idealism right: what I perceive is not any material object but rather objects that are constituted by a series of interrelated mental events, some of which I directly perceive? Although I am confident that idealism is false because it lacks any reasonable explanation for why I have the experiences I have, I am quite unsure whether direct or

indirect realism represents the stronger theory of perception. Meanwhile, I have a very high degree of justified confidence that I perceive my laptop, hands, and lap as I type away. Sometimes, we know, or very reasonably believe, certain judgments without having as much confidence in any theory from which those judgments might be derived or receive support. This is the case in ethics.

The insight that some ethical judgments more specific than an ethical theory are more credible than any theory helps to motivate the coherence model of ethical justification. To explain this idea, it will help to revisit the application model of ethical justification. Again, according to this model, we should arrive at justified ethical judgments about concrete issues by applying the correct or most justified ethical theory. Note that this method makes sense only if we are more justified in believing some ethical theory than we are justified in believing more specific ethical judgments. As discussed earlier, though, there is genuine uncertainty about which ethical theory is best. And, as just noted, we have more justified confidence in some more specific norms than we have in any theory. These insights are consistent with the alternative offered by the coherence model.

According to this model, *there is no epistemically privileged level of ethical discourse*. That is to say, neither a theory nor a set of mid-level principles, a system of rules, an array of specific judgments, or any other structure of norms would afford the sort of rock-solid certainty that would vindicate an exclusive foundation for ethics.<sup>22</sup> Where we have justified confidence in a norm—where it seems beyond reasonable doubt and our confidence does not appear to be the result of self-interested bias or other factors known to make ethical judgment unreliable—we may think of the norm as a *considered judgment*. Such a judgment is a provisional fixed point in our ethical reasoning: we take it as correct (while remaining open, at least in principle, to the possibility that some consideration will later cause us to question it). *Considered judgments are available at multiple levels of generality*. They include the principle that intentionally harming others tends to be wrong (non-maleficence), rules that prohibit slavery and sexual assault, and the judgment that that woman's kindness to the lost child was admirable and right. A good ethical theory must be consistent with most of our considered judgments. It should also account for them, as much as possible, in terms of more general norms. Our specific judgment that the woman's kind assistance to the child was admirable and right, for example, can be explained in terms of the *pro tanto* rightness of acting beneficently, the virtue of kindness, and details about the case that make those values salient.

John Rawls popularized this model of ethical justification, which he called *reflective equilibrium*, in the early 1970s.<sup>23</sup> Rawls did not invent this model, however. One can find similar ideas in the earlier work of his colleague W. V. O. Quine, although Quine was considering epistemology, in general, not a methodology in ethics in particular.<sup>24</sup> If I am not mistaken, Henry Sidgwick used the same method in his masterful *The Methods of Ethics*.<sup>25</sup> Probably many other predecessors to Rawls did as well. Among moral philosophers who explicitly discuss their methods, the coherence model might be the most popular in recent decades.<sup>26</sup> In order to clarify how my model of ethical argument—leveraging a sturdy norm—relates to the coherence model of justification, I need to elaborate on the latter model more fully.

In this model, a particular moral judgment is justified if it is a member of the most coherent possible set of moral norms. The term “coherent” is used here in a very broad sense that should be taken to incorporate all “ethical theoretical virtues”—that is, all of the characteristics that we reasonably seek in an ethical theory. One way to characterize coherence in this broad sense is in terms of the best possible realization of these features (1) logical consistency, (2) argumentative support (support of less-than-obvious norms with arguments rather than mere assertion), (3) plausibility (being believable independently of support from other considerations), (4) global illumination or comprehensiveness (covering more areas of ethics rather than fewer), (5) consistency with empirical evidence and the best empirical theories, and (6) simplicity (minimizing the number of basic ethical terms without excessive sacrifice of other theoretical virtues).<sup>27</sup>

Presumably, no one has ever arrived at an ethical theory or outlook that completely satisfies all of these aims. And it is controversial which existing ethical theory comes closest to doing so. But, the idea of full coherence is a sort of ideal toward which good ethical reasoning may be seen as working, making progress over time.<sup>28</sup> Appreciating this point permits us to avoid paralysis in trying to use this model and clarifies its compatibility with leveraging a sturdy norm.

Consider a misleading picture of how to use the coherence model of justification. According to this picture, we assume that an ethical judgment is justified if and only if it is part of a maximally coherent set of ethical judgments. Then, we try to determine whether a judgment we are evaluating—say, that a just society would not permit college students to incur crippling debt—is part of the maximally coherent set. But wait! We do not know what the perfectly coherent set of judgments is. If we did, we would already know how to answer our question. But, without knowing what the perfectly coherent set of judgments is, we cannot know whether our judgment is part of that set. Hence paralysis.

Fortunately, there is a more helpful picture of how to think about justification in terms of coherence. Consider again whether justice requires preventing college students from incurring crippling debt. We ask whether an affirmative answer would contribute more to the coherence of the overall set of ethical norms that we accept, or are inclined to accept, upon reflection than a negative answer would contribute. But how can one know? It seems humanly impossible to reflect on all the ethical judgments one believes to be true, a massive set that is (hopefully) coherent. So how to proceed?

*It is best, I suggest, to regard overall coherence as a deliberative ideal, to keep it in the background as a possible checkpoint, and to think more “locally” about ethics.* When we consider a particular ethical issue, such as the one about justice and student debt, we think about what it would mean to accept an affirmative answer and what it would mean to accept a negative answer. We look for “ripple effects”—discernible implications of each possible answer and how they relate to other ethical beliefs we hold. In my own case, if I think of the judgment that justice does *not* require policies that prevent crippling debt for college students, I note this judgment’s consistency with my belief that promises should generally be kept and that agreeing to pay back a loan is a kind of promise. It is also consistent with my belief that we should not encourage people to spend money irresponsibly—say, by blowing it on inessential goods because one knows one will be excused from most of one’s student loan. At the same time, I note that the aforementioned judgment conflicts with my belief that a wealthy society should treat education, including higher education, as primarily a social investment. It also conflicts with my belief that Americans should strive for equal opportunity to access high-quality schooling in the sense that merit, not economic status, should be the main factor in determining admission to universities. I suspect that, if I pursued this issue more thoroughly, I would find that an affirmative answer—that our society *should* have policies that prevent crippling university student debt—would prove to fit in, or cohere, better with most or all of the ethical judgments I believe to be correct. Importantly, I do not have to consult all of my ethical beliefs—a seemingly impossible task. Rather, I need to reason locally, looking for “ripple effects” on relevant beliefs that are closely related to the issue at hand.

Here is a suggestion about the relationship between leveraging a sturdy norm and the coherence model of ethical justification: *Every good ethical argument can be understood as implicitly appealing to at least one theoretical virtue embedded in the broad sense of “coherence.”* The explanation for this connection is that good ethical argumentation is not merely effective at persuading others (as manipulative appeals to emotions might also be); it is also of high quality, which requires that it be logically consistent, consistent with available evidence, and in keeping with other ethical theoretical virtues. Thus, good ethical argumentation is consistent with the coherence model of ethical justification. Moreover, the former may be understood as implicitly appealing to the theoretical virtues incorporated into the latter model. In addition, good ethical argumentation may contribute to ethical theory construction by helping to show that some ethical theory is more coherent—in the stipulated broad sense of “coherent”—than its competitors.

For an example that illustrates the italicized claim in the preceding paragraph, consider Thomson’s two arguments by analogy. Each calls for us to be consistent in our thinking about the permissibility of declining to maintain a stranger’s life at considerable cost to ourselves and about the ethics of abortion. An alternative way of thinking about these two arguments is that they demand we either provide argumentative support for any claim that it is permissible to detach from the violinist or evict the people-seed yet impermissible to abort an unwanted pregnancy or, in the absence of cogent argumentative support, acknowledge the permissibility of abortion. Yet, a third possible construal is that, if one (somewhat unusually) claims that both abortion and either detaching the violinist or evicting the



people-seed are impermissible, one may reply that it is intuitively implausible that detaching or evicting in these imagined scenarios would be wrong.

It should not surprise us that there are multiple ways to understand how a good ethical argument works in terms of engaging theoretical virtues: coherence itself is a holistic concept that applies to a whole set of norms rather than to one in particular. At the same time, an argument might highlight one theoretical virtue, one that is thought likely to be absent in thinking about the issue at hand. For example, Woolhandler and Himmelstein's case for national health insurance emphasized factual illumination about the American and Canadian healthcare systems because the authors believed that mistaken factual assumptions fueled American bias against this system of healthcare finance.

### Concluding Reflections

Over the years, I have taught many sections of a class called "Ethics: Theory and Applications," half of which is devoted to ethical theory and half to issues in applied ethics. Students sometimes express surprise that they are able to do what appears to be solid work on the applied ethical issues without being sure which ethical theory is strongest. Both in teaching and conducting research I have sometimes marveled at how few of the most powerful ethical arguments I have encountered appeal to an ethical theory. If this paper's analysis has been correct, we have a very clear explanation for why high-quality and effective ethical argumentation can proceed independently of any ethical theory: the structure of a good ethical argument is the leveraging of a sturdy norm, which need not be (and usually is not) as general as a theory.

Our reflections also clarify why those ethicists who embrace the coherence model of ethical justification so rarely appeal to it explicitly. In most cases, good ethical arguments need to appeal, only implicitly, to one or more criteria of coherence. They need not ask one to consider the entire body of one's ethical commitments and whether a particular judgment coheres with the entire structure. Instead, in doing ethics, we may think more "locally." If the issue is whether we should have policies that allow college students to incur crippling debt, we should compare the "ripple effects" of an affirmative answer to those of a negative answer and ask which answer, *as far as we can see*, results in a greater realization of the theoretical virtues embodied in our broad sense of coherence. This task is manageable. It is also so natural that, when we perform it, we ordinarily do not even realize we are doing so.

**Competing interest.** The author declares no competing interests exist.

### Notes

1. The author would like to thank members of the National Institutes of Health Department of Bioethics for helpful feedback on a draft of this paper that was presented on July 25, 2023.
2. An ethical theory might comprise a single principle, as with utilitarianism and Kant's ethics (to be discussed), or it might consist in a plurality of principles, rules, or other norms if the theorist holds that the plural norms cannot be reduced to simpler and more general terms. For an ethical theory comprising two irreducible principles, see Frankena W. *Ethics*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall; 1973: Chap. 3. For an ethical theory consisting in multiple rules and a decision procedure for addressing conflicts among them, see Gert B. *Morality: A New Justification of the Moral Rules*. New York, NY: Oxford University Press; 1988.
3. For a classic representative of this theory, see Bentham J. *The Collected Works of Jeremy Bentham: An Introduction to the Principles of Morals and Legislation*. Burns JH and Hart HLA, eds. Oxford, UK: Clarendon; 1996 (first published 1789).
4. Fletcher J. *Humanhood: Essays in Biomedical Ethics*. Buffalo, NY: Prometheus; 1979: Chap. 12.
5. I claim no originality in enumerating these difficulties, which are well-known to those working in practical ethics.

6. For a seminal discussion of this issue, see Parfit D. *Reasons and Persons*. Oxford, UK: Clarendon; 1984: Appendix I.
7. See, e.g., Hare RM. *Moral Thinking*. Oxford, UK: Clarendon; 1984.
8. See, e.g., Hooker B. *Ideal Code, Real World*. Oxford, UK: Oxford University Press; 2000.
9. See, e.g., Ross WD. *The Right and the Good*. Oxford, UK: Oxford University Press; 1930 and Nozick R. *Anarchy, State, and Utopia*. New York: Basic Books; 1974.
10. For seminal contributions to ancient and contemporary virtue ethics in the Western tradition, see, respectively, Aristotle. *Nicomachean Ethics*. Translated by Ross WD. Oxford, UK: Oxford University Press; 1956 (written in 4th century BCE) and McIntyre A. *After Virtue*, 2nd ed. Indianapolis, IN: University of Notre Dame Press; 1984.
11. See, especially, Gilligan C. *In a Different Voice*. Cambridge, MA: Harvard University Press; 1982, although this is primarily a work in moral psychology. For a more explicitly normative contribution to the ethics of care, see, e.g., Tong R. The ethics of care: A feminist virtue ethics for healthcare practitioners. *Journal of Medicine and Philosophy* 1988;23:131–52.
12. See Nussbaum M. *Women and Human Development: The Capabilities Approach*. Cambridge, UK: Cambridge University Press; 2000.
13. In other words, she must consider the matter of rhetoric, as Aristotle defined it—“the faculty of observing in any given case the available means of persuasion” (Aristotle. *The Rhetoric*. Translated by Roberts WR. In: Corbett E. *The Rhetoric and the Poetics of Aristotle*. New York: The Modern Library; 1984 [written in 4th century BCE]: Bk. I, Chap. 2)—and, assuming she does not try to persuade simply by stirring emotions or through her own credibility, she must engage in the third kind of rhetoric in Aristotle’s taxonomy, persuasion through argument (ibid).
14. For devastating critiques of both of these theories, see Kymlicka W. *Contemporary Political Philosophy*, 2nd ed. Oxford, UK: Oxford University Press; 2002: Chaps. 4 and 5.
15. Thomson JJ. A defense of abortion. *Philosophy and Public Affairs* 1971;1:47–66.
16. Singer P. Famine, affluence, and morality. *Philosophy and Public Affairs* 1972;1:229–43. Because Singer motivates this principle with a thought-experiment about saving a child who has fallen into a pond, one could construe his argumentative strategy as involving argument by analogy: In deciding whether to help the global needy with financial contributions, we are in a relevantly similar position to the person who sees a child at risk of drowning and knows that she can save the child at only modest cost to herself. For another argument falling under the strategy of appealing to an intuitively plausible norm, see DeGrazia D. Moral vegetarianism from a very broad basis. *Journal of Moral Philosophy* 2009;6:143–65.
17. Regan T. *The Case for Animal Rights*. Berkeley, CA: University of California Press; 1983.
18. Regan T. The case for animal rights. In: Fox MW and Mickley LD, eds. *Advances in Animal Welfare Science* 1986/87. Washington, DC: The Human Society; 1986:179–89.
19. For now I am relying on an intuitive grasp of the idea of a considered judgment. The concept will be explored more fully in a later section.
20. Macklin R. Ethical relativism in a multicultural society. *Kennedy Institute of Ethics Journal* 1998;8:1–22. This article was reprinted in DeGrazia D, Mappes T, and Brand-Ballard J, eds. *Biomedical Ethics*, 7th ed. New York, NY: McGraw-Hill; 2011:132–41.
21. Woolhandler S and Himmelstein D. A national health program: Northern light at the end of the tunnel. *JAMA* 1989;262:2136–37.
22. A parallel point, I suggest, applies to epistemology in general: there is no infallible foundation from which all further knowledge or justified belief can be derived. If I am right, then Descartes was on the wrong track in trying to identify a foundation that was beyond all possible doubt and accepting nothing less as a basis for knowledge. See Descartes R. *Meditations on First Philosophy*, 3rd ed. Translated by Cress D. Indianapolis, IN: Hackett; 1993 (first published 1641). Unfortunately, much of modern philosophy followed Descartes’ lead and ended up in the absurd position of motivating doubt about the existence of an external world.
23. Rawls J. *A Theory of Justice*. Cambridge, MA: Harvard University Press; 1971. The odd name stems from the idea that one should seek as much equilibrium as possible, upon reflection, between a

promising ethical theory and one's considered judgments, sometimes adjusting the theory to fit considered judgments, sometimes rejecting the latter if they do not square with the theory.

24. See, e.g., Quine WVO. Two dogmas of empiricism. *Philosophical Review* 1951;60:20–43. Interestingly, one of the blurbs on the back cover of Rawls' book is from Quine.
25. Sidgwick H. *The Methods of Ethics*, 7th ed. London, UK: Macmillan; 1907. For a good discussion, see Richardson H. Commensurability as a prerequisite of rational choice: An examination of Sidgwick's position. *History of Philosophy Quarterly* 1991;8:181–97.
26. For examples of scholars who endorse it, see [note 17](#), Regan T. 1983:135; Kagan S. *The Limits of Morality*. Oxford, UK: Clarendon; 1989:11–5; Richardson H. Specifying norms as a way to resolve concrete ethical problems. *Philosophy and Public Affairs* 1990;19:279–310, esp. Sect. IV; and Beauchamp TL and Childress JF. *Principles of Biomedical Ethics*, 8th ed. New York: Oxford University Press; 2019: Chap. 10.
27. For a fuller discussion, see DeGrazia D and Millum J. *A Theory of Bioethics*. Cambridge: Cambridge University Press; 2021: Chap. 2.
28. I am attracted to a thesis Derek Parfit advanced late in his life: that the best versions of the most promising ethical traditions were gravitating toward a single ethical theory, just as several individuals might start in very different places at the bottom of a mountain and later meet at its summit. See Parfit D. *On What Matters*, Vol. 1. New York, NY: Oxford University Press; 2011.