The fundamental divisions in ethics

Matthew Hammerton

Routledge
Taylor & Francis Group

Check for updates

The fundamental divisions in ethics

Matthew Hammerton

School of Social Sciences, Singapore Management University, Singapore, Singapore

ABSTRACT
What are the fundamental divisions in ethics? Which divisions capture the most important and basic options in moral theorizing? In this article, I reject the 'Textbook View' which takes the tripartite division between consequentialism, deontology, and virtue ethics to be fundamental. Instead, I suggest that moral theories are fundamentally divided into three independent divisions, which I call the neutral/relative division, the normative priority division, and the maximizing division. I argue that this account of the fundamental divisions of ethics better captures the main concerns that normative ethicists have when assessing moral theories. It also helps us make progress in comparative ethics and makes visible theoretical possibilities obscured by the Textbook View.

1. Introduction

Moral philosophers construct theories that attempt to systematize our moral thinking. We compare these theories to each other and debate which theories give the best account of morality. Sometimes this comparison concerns very specific theories. For example, we might compare W.D. Ross's 'ethics of prima facie duties' with G.E. Moore's 'ideal utilitarianism' and assess their respective merits. However, often the focus is more general. Instead of looking at specific theories, we look for broad theoretical categories that all moral theories fall under and debate whether the best approach to ethics would fall under one category or another. When we do this we are attempting to capture the most fundamental divisions in ethics – the distinctions that are most relevant to the task of systematizing our moral thinking. This is the topic that this article is concerned with. I will ask: which divisions are most important

CONTACT Matthew Hammerton ✉ mhammerton@smu.edu.sg 🏛 School of Social Sciences, Singapore Management University, 90 Stamford Road, Singapore 178903, Singapore
This article has been corrected with minor changes. These changes do not impact the academic content of the article.

to moral theorizing, capturing the main options available when construct-ing a moral theory?

A standard answer to this question divides moral theories into three general categories – *consequentialism*, *deontology*, and *virtue ethics*. I will call this answer the 'Textbook View' because it often appears in intro-ductory texts. The Textbook View presents its three categories as the three main approaches to ethics, marking the most important differences that can exist between moral theories. Often, these three categories are understood as mutually exclusive and exhaustive.

It is fair to say that the Textbook View, in some form or another, is widely endorsed. Its three categories are treated as fundamental not only in ethics textbooks but also in many academic works by contempor-ary moral philosophers.[1] Furthermore, even those who think that it is a little too simplistic often agree that it is a useful approximation of our nor-mative ethical options. And, those who have objected to aspects of the Textbook View usually present their challenge as a modification to it rather than as a reason for rejecting it entirely.[2]

In this article, I argue that the Textbook View is deeply flawed and needs to be replaced with an alternative picture of the fundamental div-isions in moral theory. The view I favour locates moral theories along three major independent divisions, which I call the *neutral/relative div-ision*, the *normative priority division*, and the *maximizing division*. These three divisions can be contrasted with the single tripartite division offered by the Textbook View. To make the case for this alternative picture I review each of its three divisions in §3 and explain how each cap-tures different areas of concern that we have in our moral theorizing. In §4, I then compare these divisions to the Textbook View and demonstrate that they are incompatible. I conclude in §5 by considering some benefits of adopting the alternative view I offer in place of the Textbook View.

2. What are 'fundamental' divisions?

Any account of the *fundamental divisions in ethics* presupposes that some divisions can be meaningfully described as 'fundamental'. Therefore, a

---

[1]Several examples are cited in §§4.1–4.3.
[2]For example, some believe that the Textbook View is not exhaustive because there are moral theories that do not fit under any of its three categories. Moral particularism, care ethics, and the 'hybrid theory' of Scheffler (1994) are the most commonly cited examples here. Yet, those who take this line generally suggest modifying the Textbook View to include these extra categories rather than abandoning it. The arguments I make against the Textbook view also apply to various 'Textbook+' views that add a handful of additional categories.

helpful starting point is to clarify what the term 'fundamental' means in this context. I endorse the following account offered by Peter Vallentyne:

> To say that a distinction is fundamental is to say that it is not significantly less important than any other distinction. Fundamental distinctions are ones that most ultimately matter. The importance of a distinction, I assume, is relative to a set of interests and purposes. Here we are concerned with the importance of distinctions relative to our interest in the assessment of moral theories. The importance of a distinction relative to these interests is determined by something like the usefulness of the role it plays, or would play on reflection, in the criticism and justification of moral theories. (1987, 29)

I claim that the three divisions I set out below are superior to the Textbook View because they better capture the things that we are interested in when we assess moral theories. In contemporary moral philosophy done in the analytic tradition, there is a broad consensus on what kinds of factors are relevant to the assessment of moral theories. For example, it is widely agreed that the acceptability of the deontic verdicts a moral theory produces is relevant to its assessment. For each of the divisions I present below, I will explain how it accurately captures certain key considerations we have when assessing moral theories. Later I will show that the Textbook View inadequately captures these considerations even though it is often framed in terms of them.

An interesting question to pose here is whether the considerations I appeal to, which are widely accepted in contemporary analytic moral philosophy, should be applied to the assessment of moral theories from other philosophical traditions. I will not attempt to answer this question in this article. However, I will touch on it in §5 when I consider how certain debates in comparative moral philosophy could be pursued more fruitfully if framed in terms of the divisions I defend rather than the Textbook View.

3. Three fundamental divisions in ethics

3.1 The neutral/relative division

Some moral rules, reasons, and theories give agents a special focus on themselves, their actions, or states of affairs connected to them. Other rules, reasons, and theories, give agents a general concerns with actions or states of affairs that may involve anyone. Moral philosophers generally label rules, reasons, and theories of the former kind as 'agent-

relative' and those of the latter kind as 'agent-neutral'.[3] Several accounts of the agent-relative/agent-neutral distinction have been defended. Parfit's (1984) account is perhaps the most intuitive. He holds that all moral rules give agents certain substantive aims. For example, a rule prohibiting lying gives each agent the substantive aim that she does not tell lies. Parfit then defines agent-neutral rules as rules that give all agents the same aims and agent-relative rules as rules that give different aims to different agents. For example, consider the following four rules:

(1) Each agent must not tell lies.
(2) Each agent must ensure, to the best of her ability, that her family is honest.
(3) Each agent must ensure, to the best of her ability, that everyone is honest.
(4) Each agent must minimize general violations of (1).

Rules (1) and (2) are standardly classified as agent-relative, whereas (3) and (4) are standardly classified as agent-neutral. According to Parfit, they are classified this way because (1) and (2) give different aims to different agents, whereas (3) and (4) give all agents the same aims. To see this, consider that (1) gives Cain the aim that *Cain does not tell lies* and gives Abel the different aim that *Abel does not tell lies*. Likewise, (2) gives Gandhari the aim that *her family is honest*, and Pandu the different aim that *his family is honest*. By contrast, (3) gives all agents the common aim that *everyone is honest*, and (4) gives all agents the common aim that *there are minimal violations of* (1).[4]

Parfit extends this account to moral theories. He classifies a theory as agent-neutral if it gives all agents the same ultimate aims and agent-relative if it sometimes gives different ultimate aims to different agents. It follows that containing at least one agent-relative rule is enough to make a theory agent-relative, whereas agent-neutral theories contain only agent-neutral rules.[5]

---

[3]Nagel (1970) was the first to properly recognize this distinction. He used the terms 'objective' and 'subjective'. Parfit (1984) expands on Nagel's work and adopts the alterative terminology of 'agent-neutral' and 'agent-relative'. Parfit's terminology is now widely followed, including by Nagel himself (Nagel 1986).

[4]See McNaughton and Rawling (1993, 83–86) and Hammerton (2019) for an important clarification of these standard classifications that shows how they resist a common line of criticism.

[5]One might wonder why theories are defined in this asymmetrical way. Why not say instead that: (1) agent-neutral theories contain only agent-neutral rules, (2) agent-relative theories contain only agent-relative rules, and (3) a theory that mixes neutral and relative rules is a 'hybrid theory'. The answer is that it is widely accepted that any plausible moral theory will contain at least some agent-neutral rules. What moral philosophers tend to disagree about is whether the best candidate moral theories will also contain agent-relative rules. Thus, applying the distinction to theories in

Parfit's account of the distinction is widely accepted. However, some question his assumption that all moral rules give agents aims, or doubt the reliability of his appeal to intuitive judgments about the commonality of aims.[6] These concerns can be addressed by turning to formal accounts of the distinction that mark it by pointing to differences in the logical structure of agent-relative and agent-neutral rules. In particular, the formal account developed by McNaughton and Rawling (1991) and Hammerton (2019) addresses these concerns while fitting well with Parfit's account.

Although the agent-relative/neutral distinction is the most widely known distinction of its kind, there is a family of related distinctions that work in a similar way. For example, a moral rule or theory can be time-relative or time-neutral (Parfit 1984) and world-relative or world-neutral (Dreier 2018). Time-relative rules give agents different aims at different times and world-relative rules give agents different aims in different possible worlds.[7]

The agent-neutral/agent-relative distinction (and the other related distinctions) are widely regarded as an important development in our understanding of morality.[8] There are two main reasons why the neutral/relative divisions are important for moral theorizing. First, relative theories are able to produce certain commonsense deontic verdicts that neutral theories cannot produce. Most notable here is that agent-relative theories are compatible with deontic constraints and special duties whereas agent-neutral theories are not.[9] Deontic constraints prohibit agents from performing certain acts even when doing so is the only way to prevent more acts of that type being performed by others. Special duties require agents to prioritize those whom they have a special relationship with even if not prioritizing such people would result in more special-relationship prioritization overall. Both are core commitments of our commonsense morality. In addition to these well-known examples, there are also certain deontic verdicts that can only be produced by a time-relative, or a world-relative theory.[10] Because we give

---

this asymmetrical way has the virtue of capturing the distinction that is of most interest to moral philosophers.

[6]For example, see Schroeder (2011, 36–39), and Portmore (2013).

[7]See Dreier (2018) and Hammerton (2020) for further discussion.

[8]Among those who emphasize its importance are Nagel (1970, 1986), Parfit (1984), McNaughton and Rawling (1991), Dreier (1993), Hurka (2003), and Portmore (2013).

[9]For further discussion of this point see Dreier (1993), Dougherty (2013), Nair (2014), Hammerton (2017), and Cox and Hammerton (2021).

[10]See Parfit (1984), Broome (1991), Dreier (2018), and Hammerton (2020) for further discussion.

weight to how well a moral theory captures commonsense morality, a theory's status as relative or neutral is one of its key features.

The second reason why the neutral/relative division matters is that neutral theories give all agents, at all times and places a common moral outlook, whereas relative theories give different moral outlooks depending on which agent you are, what your temporal location is, and what possible world you inhabit. Several have argued that moral theories are more or less plausible depending on whether they give a common moral outlook to everyone everywhere or give different moral outlooks to agents depending on their position. Sen (1982) argues that giving different moral outlooks is more plausible because who we are, and where and when we are located, are morally significant factors that are relevant to how we ought to evaluate and respond to the world. By contrast, Parfit (1984) argues that giving different moral outlooks to agents at different positions is problematic because it results in Prisoner Dilemma scenarios where a moral theory is 'directly collectively self-defeating'. Relatedly, Pettit (1997) argues that it is problematic because it leads to a kind of 'moral civil war' between agents, or even different time-slices of the same agent, as each pursues their different moral ends. Given these points, whether a theory is neutral or relative appears to be a key thing to consider when assessing its overall plausibility.

3.2. Normative priority

Moral discourse contains several distinct normative concepts. A list of such concepts might include good, bad, right, wrong, just, unjust, virtue, vice, praiseworthy, blameworthy, and reason. Each of these concepts purports to pick out a kind of moral property. For example, rightness is a property of actions, goodness a property of states of affairs, and virtue a property of character traits. When discussing these properties, moral philosophers sometimes talk about one of them 'being prior' to another. For example, influential works by Ross (1930), Broad (1930), and Rawls (1971) prominently discuss whether 'the good is prior to the right'.

What is meant by this priority talk is often left obscure. Sometimes it makes a conceptual claim, where 'A is prior to B' means that the B-concept is defined in terms of the A-concept. Sometimes it makes a metaphysical claim, where 'A is prior to B' means that an appropriate asymmetric metaphysical relation holds between A-properties and B-

properties. Options for this relation include reduction, supervenience, grounding, metaphysical analysis, metaphysical priority, and ontological dependence. It may even be used to make several of these claims at once. It is not necessary to settle in this article which of these interpretations is best. As I argue below, the main function of priority talk in normative ethics is to show how moral theories explain the normative claims they make. Furthermore, all the interpretations of priority talk considered above entail an explanatory relationship between A and B.[11] Therefore, henceforth I will interpret 'A is prior to B' as 'A-properties explain B-properties'.

In addition to priority talk, it is also useful to have the notion of a *normative primitive*. A normative primitive is a moral property that is not explained by any other *moral* properties. They are the most basic moral properties in a normative system. This says nothing about their standing more generally and thus leaves open the question of whether any particular normative primitive can be given a non-moral explanation.

Understanding priority and normative primitives allow us to clarify the various explanatory relations that can hold between different moral properties. For any two moral properties, A and B, one of the following is true:

  (i) A explains B.
 (ii) B explains A.
(iii) A and B are explained by other moral properties.
(iv) A and B are each normative primitives.
 (v) One of A and B is a normative primitive, the other is explained by another moral property.
(vi) A and B are actually one and the same property.

As an example, suppose that A is the property of *having a reason to desire that P*, and B is the property of being a *good* state of affairs. Perhaps states are good because we have a reason to desire that they obtain (option i). Or perhaps it is the goodness of certain states that explains why we have a reason to desire that those states obtain (option ii). Or maybe the goodness of a state S, and our reason to desire S, are both explained by S being something that a virtuous agent would desire (option iii). Or perhaps S being good, and there being a reason to desire S, are basic normative facts (option iv). Or,

---

[11]With the exception of the conceptual interpretation which only entails that, if there are A-properties and B-properties, then the former explains the latter. However, this is good enough.

maybe S being good is a basic normative fact, whereas our reason to desire S is explained by S being something that a virtuous agent would desire (option v). Or, finally, maybe 'S is good' and 'there is a reason to desire that S' are different ways of describing the same normative property (option vi).

Articulating the relations that occur between all the moral properties in a moral theory gives us a complete picture of the internal explanatory structure of that theory. In fact, we can sketch the structure of several major moral theories and compare them (Figures 1–15). To do this I will use an arrow to mark normative priority (the arrow will point from the prior to the non-prior). A dashed arrow will mark a partial connection. For example, a theory might hold that A-properties are sometimes explained by B-properties and sometimes explained in other ways. Finally, normative primitives are written in uppercase, non-primitives in lower-case.

$$\text{VALUE} \longrightarrow \text{rightness}$$

Figure 1. *Simple teleology*. The classical utilitarian theories of Bentham ([1789] 1907) and Mill ([1861] 1998) are widely interpreted this way. Various theories standardly regarded as 'non-consequentialist' also have this structure. For example, standard natural law theory, as defended by Finnis 1980, has this structure, deriving moral duties from what would *honour* various final goods. Also, Francis Kamm's moral theory has this structure. Kamm (2007, 29) says: ' … my account highlights an agent-neutral value: the high degree of inviolability of persons that lies at the base even of an agent-relative duty'. Finally, some Kantians, such as Korsgaard (1996), appear to attribute this structure to Kant's moral theory.

$$\text{GOOD FOR} \longrightarrow \text{rightness}$$

Figure 2. *Ethical egoism*. See Sidgwick (1907) for a classic account of ethical egoism and Burgess-Jackson (2013) for a recent sympathetic discussion of it.

$$\text{VALUE} \rightrightarrows \genfrac{}{}{0pt}{}{\text{rightness*}}{\text{virtue}}$$

Figure 3. *Global consequentialism*. Global consequentialism is suggested by Parfit (1984, 24–28) and defended by Kagan (2000) and Pettit and Smith (2000). Note that 'rightness' is starred because, unlike in other diagrams where it only applies to actions, here it applies to all categories (e.g. institutions, motives, feelings, etc.).

$$\text{FITTINGNESS} \longrightarrow \text{value} \longrightarrow \text{rightness}$$

Figure 4. *Fittingness teleology*. Recent defences of fittingness teleology include Garcia (1986), Smith (2003), Chappell (2012), and Cullity (2015).
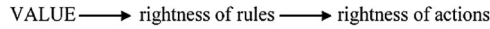
VALUE ⟶ rightness of rules ⟶ rightness of actions

Figure 5. *Rule consequentialism*. See, Brant (1979) and Hooker (2000) for book length defences of rule consequentialism.

RIGHTNESS ⟶ value
RIGHTNESS ⟶ virtue

Figure 6. *Rightness Priority*. Rawls (1971, 491) is the most clear-cut example of someone endorsing a moral theory with this structure. He says: 'We should therefore reverse the relation between the right and the good proposed by teleological doctrines and view the right as prior. The moral theory is then developed by working in the opposite direction'. Vallentyne (1987, 24) suggests that Kantianism is best interpreted as having this structure. Some Kantians, such as Cummiskey (1996), and Korsgaard (1996) explicitly interpret Kantianism as not having this structure, whereas others such as Herman (1993) give accounts that are compatible with Kantianism having this structure without explicitly stating that it does. Frankena (1973, 48) is a classic presentation of the view that virtue is derived from right action.

VALUE ⤏ prima facie duty
virtue ⟶ rightness
PRIMA FACIE DUTY ⟶

Figure 7. *Rossian intuitionism*. According to Ross (1930), moral rightness is explained by the five prima facie duties and their respective weights. Some of these prima facie duties are in turn directly explained by the value of states of affairs (which is a normative primitive for Ross) or by virtuous character traits (which are ultimately explained by value). However, other prima facie duties cannot be given further explanation and are themselves normative primitives. Because the connection between prima facie duties and value is partial, a dashed arrow is used. By contrast, right actions are always explained by appeal to all of the prima facie duties, hence the non-dashed arrow.

REASON ⟶ rightness
REASON ⟶ value

Figure 8. *The buck-passing view of value*. This view is defended by Scanlon (1998), who gives it its name. It is also discussed in Broad (1930).

EUDAIMONIA ⟶ virtue ⟶ rightness

Figure 9. *Standard neo-Aristotelian ethics*. Several moral theories inspired by Aristotle, such as Hursthouse (1999) and LeBar (2013), have this structure. See Crisp (2015) and Hirji (2018) for arguments that, contrary to what is commonly thought, this is not the structure of Aristotle's moral theory.

ATTRIBUTIVE GOOD → virtue
→ rightness

Figure 10. *Footian virtue ethics*. This view is defended by Foot (2001). She uses the term 'human goodness' to refer to the attributive use of 'good' applied to humans.

ADMIRABILITY → virtue → rightness
→ value

Figure 11. *Agent-based virtue ethics*. This view is developed and defended by Slote (1992) and (2001). Note that the normative primitive here is admirability applied to motives *only* and not to any other category.

VALUE → virtue → rightness

Figure 12. *Neo-Platonist virtue ethics*. Murdoch (1971) and Chappell (2014) defend versions of this view.

virtue
↑
VALUE → rightness

Figure 13. *The recursive account of virtue*. This theory is developed by Hurka (2001). He posits a set of basic goods and recursively defines virtues as additional goods that consist of loving the basic goods. He then suggests a 'consequentialist' view according to which right action involves promoting the good (where this includes both the basic goods and recursively defined goods [i.e. virtues]).

VIRTUE ⤸
↕   RIGHTNESS
VALUE ⤸

Figure 14. *Holistic virtue ethics*. According to this view, virtue, rightness, and value all partially depend on each other. Annas (1993, 7–10) attributes this view to several ancient philosophers. Annas (2011, Chapter 10) directly defends it herself.

VALUE - - - - -> reasons
VIRTUE - - - - - - - - - - -> rightness
REASONS - - - - - - - - - ->

Figure 15. *Moral particularism*. Dancy (1993, 2004) holds that all act tokens are right or wrong for particular reasons. However, because of his reasons holism, he thinks that there are no exceptionless generalizations about what these reasons are. He also appears to accept pluralism about the kinds of factors that are relevant. Sometimes, facts about value or virtue are part of the explanation of the reasons for action in a specific case. At other times, these reasons are normative primitives not explained by further moral facts. The normative structure of Dancy's position closely resembles that of Ross. The main visible difference is that, whereas Ross thinks that all the prima facie duties explain rightness, Dancy thinks different reasons will often explain the rightness of different specific acts. Also, in Ross's theory, the partial connections (picked out by the dotted lines) are few in number and ruled governed. By contrast, in Dancy's theory, there are many of them (because of the focus on act tokens) and they are not rule-governed.

Diagrams of further moral theories could be added here. What is striking as we look through these diagrams is the sheer variety of normative structures available. Rawls (1971) famously asks whether the good is prior to the right or the right prior to the good, and treats this as the central distinction when it comes to a moral theory's normative structure. However, the fifty years of normative ethics research output that has followed Rawls's canonical work demonstrates that the menu of options is much richer.[12] When we think about normative priority, it is important to keep this diversity in mind. The fundamental division we find here is not a simple bipartite one but a multi-part one that divides and sorts all moral theories by their normative structure.[13]

There is one key reason why moral philosophers are interested in normative priority. The normative structure of a moral theory tells you nothing about the deontic, evaluative, or aretaic verdicts it endorses.[14] However, we are interested in more than just the verdicts a moral theory produces. We are also interested in whether the theory gives a plausible explanation of why these verdicts hold true. Different normative structures correspond to different explanatory stories. Some of these explanations are more

---

[12]In fact, more extensive 'menus' were already on offer prior to Rawls. For example, see the influential taxonomy of Broad (1930), which includes many, but not all, of the options suggested here.

[13]I put aside the question of how this sorting is best done, although this question will receive some discussion in §4 below. Clearly, it can be done in more fine-grained or coarse-grained ways. The key point for now is that, when dividing moral theories by their most general features, one thing to focus on is the different normative structures they have. We can recognize this without taking a stand on the exact number of basic categories that divide normative structures.

[14]This point is further elaborated on in §4. One exception is when a theory leaves out certain moral properties because it holds that there are no true verdicts related to that property.

compelling, others less so. Therefore, classifying theories by their normative structure is crucial to assessing their overall plausibility.

A second reason for interest in normative priority is that knowing which moral properties are normative primitives is important for several meta-ethical questions. For example, if we are assessing the viability of a natural-istic reduction of ethics then it is helpful to know which moral properties are, or could be, normative primitives as these are the properties that might be directly reduced to natural properties.[15] These considerations show the importance of normative priority in the assessment of moral theories.

3.3. Maximizing the good

Recent literature that compares consequentialist and non-consequential-ist approaches to ethics has converged on an account of what the key feature of consequentialism is that makes it compelling to so many.[16] Fol-lowing Foot (1985) this feature has come to be called the 'compelling idea' of consequentialism. In its standard formulation, it is the idea that it is always permissible to do what will lead to the best outcome. There is said to be something intuitively compelling about this idea. How can you be faulted if you have done what will result in the best outcome? The Compelling Idea is closely connected to the thought that there is something deeply appealing about maximizing rationality. As Samuel Scheffler puts it:

> The kind of rationality that consequentialism seems so clearly to embody, and which makes so much trouble for views that incorporate agent-centered restric-tions, is what we may call maximizing rationality.[…] The core of this con-ception of rationality is the idea that if one accepts the desirability of a certain goal being achieved, and if one had the choice between two options, one of which is certain to accomplish the goal better than the other, then it is, ceteris paribus, rational to choose the former over the latter.[17]

Scheffler is careful here not to tie the appeal of maximizing rationality to moral goodness. Instead, he connects it to the desirability of certain ends, which could be cashed out in terms of the moral goodness instantiated in those ends or in some other account of their desirability. Therefore, the Compelling Idea can be stated more generally as the idea that it is

---

[15]See Smith (2005, 12) for further discussion.
[16]See, Scheffler (1994, 141–143), Foot (1985, 198), Portmore (2005, 95), Schroeder (2007, 279), and Dreier (2011, 100).
[17]Scheffler (1985, 252). See similar statements in Rawls (1971, 21).

always morally permissible to do what will lead to the outcome that is most morally desirable.[18]

One important thing to notice about the Compelling Idea is that whether a theory satisfies it is not determined by it normative structure. More specifically, explaining the right in terms of the good is neither necessary nor sufficient for satisfying the Compelling Idea. It is not sufficient because various moral theories explain the right in terms of the good yet do not connect them via a 'maximizing' or 'promotion' function (examples discussed in the previous section include Finnis (1980) and Kamm (2007)). It is not necessary because a theory can hold that agents are always required (or permitted) to maximize the good without explaining this requirement in terms of the goodness of states of affairs. For example, suppose I hold that a virtuous agent would always choose to maximize pleasure and I use this to explain both why pleasure is good and why we are always morally permitted to maximize pleasure. Such a theory does not explain rightness by appeal to the good, yet it satisfies the Compelling Idea.

Noticing that the Compelling Idea cannot be reduced to facts about normative priority suggests that there is a third division in ethics that is distinct from neutrality/relativity and normative priority – the distinction between moral theories that always permit agents to do what leads to the most desirable outcome and those that sometimes prohibit agents from doing this. This division is fundamental because what side of it a moral theory falls on determines whether that theory satisfies the Compelling Idea. We can call this division the 'maximizing' division because it is linked to the intuitive appeal of maximizing rationality in ethics.

In addition to the Compelling Idea, there is another reason why the maximizing division matters in ethics. Moral theories need to give deontic verdicts in circumstances in which we are uncertain about one or more of the relevant facts. Theories that incorporate expected utility theory into their structure are able to do this well. Those that do not seem to have a problem.[19] Yet to be compatible with expected utility theory, a moral theory must contain an ordering of possible outcomes, and its deontic verdicts must correspond to the output of some kind of promotion function over this ordering. On the assumption that any plausible ordering of possible outcomes is connected to the desirability of those outcomes, a theory is compatible with expected utility theory only if it satisfies the Compelling Idea. Therefore, the maximizing division

---

[18]See Portmore (2011, 34) for an account of the Compelling Idea that fits with this broader conception.
[19]For discussion, see Jackson and Smith (2006).

also appears to mark those theories that are able to adequately deal with decision under uncertainty and those that are not.[20]

4. Trouble for the Textbook

Above I argued that there are three fundamental divisions in ethics related to neutrality/relativity, normative priority, and maximizing. Each division is logically independent of the others and each matters for different, yet important reasons. Together, they capture several of the main concerns that normative ethicists have when assessing moral theories.

Yet this model is in tension with the Textbook View, which treats the single tripartite distinction between consequentialism, deontology, and virtue ethics as the most fundamental division in ethics. An obvious way to resolve this tension is to show that the Textbook View and the three divisions are in fact compatible. It is notable that accounts of consequentialism, deontology, and virtue ethics often appeal to neutrality/relativity, normative priority, or maximizing the good as key features that divide the Textbook categories. Therefore, to assess the Textbook View we need to examine the possibility that it captures, or is equivalent to, one or more of the three fundamental divisions.

4.1. The neutral/relative division

Classical utilitarianism is an agent-neutral moral theory. Most forms of consequentialism that have followed it have also been agent-neutral. This is because consequentialists have generally assumed, by default, that the good is agent-neutral. However, in recent decades this assumption has been challenged, and it is now widely recognized that a consequentialist theory could employ an agent-relative axiology, allowing it to give agent-relative deontic verdicts.[21] Therefore, the distinction between consequentialism and non-consequentialism does not correspond to the distinction between neutral and relative moral theories.

The moral theories that are classified as forms of 'virtue ethics' are typically agent-relative theories. This is because they typically side with commonsense morality in endorsing deontic constraints and special duties (which required agent-relative rules). However, the special focus on virtue (whatever it amounts to) that characterizes virtue ethics appears

---

[20]See Portmore (2016), and Lazar (2018) for further discussion.
[21]For example, see Sen (1982), Broome (1991), Dreier (1993, 2011), Hurka (2001), Portmore (2001, 2011), Smith (2003, 2009), Louise (2004), and Hammerton (2020).

compatible with agent-neutrality. For example, a moral theory could hold that virtue is the sole normative primitive and give all agents the same aim of ensuring, to the best of their ability, that everyone achieves a virtuous life. Such a theory has the special focus on virtue that is characteristic of virtue ethics and yet is agent-neutral.

It is widely held that Deontology is necessarily agent-relative. This is because endorsing deontic constraints is commonly seen as a necessary feature of deontology, and deontic constraints appear necessarily agent-relative. Recently, Dougherty (2013) has challenged this orthodoxy by arguing that deontic constraints can also be produced by an agent-neutral theory with a non-maximizing structure. Whether this argument succeeds is disputed.[22] If the argument fails, and deontology is indeed necessarily agent-relative, then there *is* a connection between it and the neutral/relative division. However, this connection would not be enough to make the distinction between deontological and non-deontological theories correspond to the neutral/relative division. This is because 'non-deontological' theories (i.e. various forms of consequentialism and virtue ethics) come in both neutral and relative varieties. Therefore, the important features that are tracked by the neutral/relative division are not tracked by the deontological/non-deontological division.

4.2. The normative priority division

Diverse normative structures are found in deontology. Amongst the various theories categorized as 'deontology' are theories that make the good prior to the right,[23] the right prior to the good,[24] and neither prior to the other.[25] Therefore, contrary to what some have thought,[26] deontology cannot be marked by a certain position taken on normative priority.

Orthodoxy holds that the good being prior to the right is a necessary feature of consequentialism. However, some have rejected this view, holding that theories which always require agents to maximize the good count as consequentialist regardless of the explanatory relationship they posit between the right and the good.[27] In favour of the orthodox view is the fact that all paradigmatic examples of consequentialism

---

[22]See Hammerton (2017) for criticism.
[23]See Finnis (1980) and Kamm (2007).
[24]See Rawls (1971) and Vallentyne's (1987) interpretation of Kant.
[25]See Ross (1930) and Scanlon (1998).
[26]See Rawls (1971), Chappell (2012, 700), and Lazar (2018, 861).
[27]For example, see Brown (2011, 753) and Setiya (2018, 94).

have held that the good is prior to the right. In favour of the alternative view is the fact that a theory can satisfy the so-called 'Compelling Idea' of consequentialism without holding that the good is prior to the right.

Because 'consequentialism' is a philosophical term of art, there may be no fact of the matter about which of these two ways of using the term is correct. However, either way, the distinction between consequentialism and non-consequentialism does not correspond to any distinction in the type of normative priority that a theory contains. This obviously follows from the alternative view. But it also follows from the orthodox view because even if the good being prior to the right is a necessary condition for consequentialism, it is not a sufficient condition as several deontological and virtue ethical theories also hold that the good is prior to the right.[28]

In normative ethics, it is now standard to distinguish 'virtue theory' from 'virtue ethics'.[29] Virtue theory is any moral theory that includes the concept of virtue somewhere in its normative structure. Virtue ethics is the class of moral theories that not only include virtue in their normative structure but also give it a central place in that structure. It is widely agreed that 'instrumentalist' accounts of virtue (which explain virtue in terms of the tendency of certain character traits to promote the good) and 'deontological' accounts of virtue (which explain virtue in terms of the character traits possessed by those disposed to obey the correct moral rules) are excluded from the category of 'virtue ethics'. However, cashing out the vague notion of 'virtue playing a central role' in a way that captures all, and only, those theories that are standardly regarded as versions of 'virtue ethics' is exceptionally difficult. One common proposal is that holding virtue to be prior to rightness is the distinctive feature of virtue ethics.[30] However, this would wrongly include Hurka's consequentialist virtue theory and Ross's moral theory under the label 'virtue ethics', and wrongly exclude the virtue ethical theories of Philippa Foot and Julia Annas (see Figures 10 and 14 above). Another proposal is that virtue having intrinsic value is the distinctive feature of virtue ethics.[31] However, like the previous suggestion, this counts Hurka's consequentialist virtue theory and Ross's moral theory as

---

[28]Deontological examples include Finnis (1980) and Kamm (2007). Virtue ethics examples include Murdoch (1971) and Chappell (2014).
[29]This distinction was introduced by Driver (1996).
[30]For example, see Frankena (1973, 63) and Slote (1992, 89).
[31]See Crisp (2015).

versions of virtue ethics. Other more complicated suggestions have similar problems.[32]

Because of these problems, some virtue ethicists admit that the concept of virtue ethics is unavoidably vague (Hursthouse 1999, 4–8), and some even suggest that we abandon the concept altogether (Nussbaum 1999). What we can certainly say is that, to date, no account of virtue ethics has been given that captures all and only those theories generally regarded as 'virtue ethics' and corresponds to some kind of relevant difference in the normative structures that moral theories have. Thus, as things stand, 'virtue ethics' does not correspond to a distinctive stance in normative priority that separates it from consequentialism and deontology. I suspect that this is because the distinctions that are most relevant to categorizing moral theories according to their normative structures cut across traditional groupings of theories into 'virtue ethics' and 'non-virtue ethics'. However, I do not have space here to argue for this claim.

4.3. The maximizing division

It is widely accepted that consequentialist theories always permit agents to maximize the good.[33] It is in virtue of this that they satisfy the 'Compelling Idea'. It is also widely accepted that all deontological theories contain constraints that prohibit agents from maximizing the good in certain circumstances. Therefore, the maximizing division separates consequentialist theories from deontological theories.

However, neither category exclusively occupies its side of the maximizing division. If the orthodox view discussed in the previous section is correct, then there are theories that always permit agents to maximize the good, yet are not versions of consequentialism because they do not make the good prior to the right. Furthermore, Scheffler's (1994) 'hybrid theory' satisfies the Compelling Idea yet is not considered a version of consequentialism because it gives agents the option of not maximizing (or satisficing) the good in certain circumstances.[34] Likewise, various forms of virtue ethics prohibit agents from maximizing the good in certain circumstances yet are not understood as versions of deontology.

---

[32]For example, McAleer (2007) suggests that virtue ethics holds both that virtue is prior to the right and that it cannot be reduced to the 'non-attributive' good. However, this wrongly excludes neo-Platonist virtue ethics (e.g., Murdoch (1971) and Chappell (2014)).

[33]Or, at least to maximize that which we have most reason to desire (see, Portmore 2011).

[34]This is how Scheffler (1994), and the literature following him, has classified his theory.

Therefore, the distinction between consequentialism and deontology (or between each category and its complement) does not correspond to the distinction between theories that always permit agents to do what will lead to the outcome that is most morally desirable and those that do not.

For similar reasons to those discussed in §4.1, virtue ethics is not connected to the maximizing division. This is because the special focus on virtue (whatever it amounts to) that characterizes virtue ethics is compatible with always permitting agents to maximize the good. For example, a simple version of virtue ethics might hold that we are required to act as a virtuous agent would act in the relevant circumstances. Yet, perhaps a virtuous agent would always act in ways that maximize good.

4.4. The overall issue

The problem with the Textbook View is twofold. First, none of its three categories neatly map onto any of the three fundamental divisions. For example, the distinction between 'consequentialism' and 'non-consequentialism', whatever it ultimately amounts to, does not correspond to the distinction between neutral and relative theories, or any distinction in normative priority, or the distinction between theories that always permit agents to maximize the good and those that sometimes prohibit it. The same holds for deontology and virtue ethics. Second, insofar as some connections do hold between the Textbook categories and the fundamental divisions, these connections are uncertain because none of the Textbook categories is used with the same precision that is built into the three divisions. For example, some tie 'consequentialism' to the idea that the good is prior to the right whereas others do not, and there seems to be no fact of the matter about which view is correct.

The upshot is that categorizing a theory according to the Textbook View does not convey much relevant information about its key theoretical properties, yet categorizing a theory with the three fundamental divisions does. To see this, suppose that there are three moral theories, A, B, and C, which you need to evaluate yet know nothing about. If I tell you that A is a consequentialist theory, B is a deontological theory, and C is a virtue ethical theory, I have not told you much at all. You may make some educated guesses about what each theory is like, but many of those guesses could turn out to be wrong. For example, you might wrongly suppose that A is agent-neutral (when it could be an agent-relative form of consequentialism), or wrongly guess that in theories B and C the good is not prior to the right (yet B might be a natural law theory and C a neo-

Platonist version of virtue ethics). On several key criteria, you will not be sure how A, B and C compare to each other. On the other hand, if I do not categorize these theories according to the Textbook, but instead tell you whether each theory is neutral or relative, whether it always permits maximizing the good, and what kind of normative structure it has, then you will be very well informed about how the three theories compare to each other.

In response to these problems, one might attempt to defend the Textbook View by arguing that its three categories are family resemblance concepts based on the historical connections and loose similarities that hold between different moral theories. Applying this idea to consequentialism, Tom Dougherty says:

> Instead, we could think of the consequentialist tradition, for example, as beginning in full force with the classical utilitarians. The tradition includes thinkers like G. E. Moore, who held a pluralist theory of the good, so that that beauty, for example, is good as well as happiness. And the tradition ends up with contemporary theorists such as those who hold that the impersonal value of someone's happiness depends on whether she deserves it. What ties these theories together into a tradition is the impact of the thoughts of earlier theorists on later theorists, and the fact that consequently these theories have certain paradigmatic features. (2013, 536)

Similar claims could be made about 'deontology' and 'virtue ethics'. I am sympathetic to this way of understanding the Textbook categories. It may be their most charitable interpretation. However, it greatly diminishes the role they play in systematizing our normative theorizing. Historical lines of influence are not the same thing as fundamental divisions. That two theories share a historical connection and certain paradigmatic features does not entail that what is distinctive about them corresponds to a fundamental division. Furthermore, sorting moral theories by their historical lines of influence is rather parochial. There might be an historical line of influence between Aristotle's moral theory and Michael Slote's agent-based virtue ethics. However, there is no historical line of influence between these theories and Confucianism. Hence, asking whether Confucianism belongs to the 'virtue ethics tradition' does not make sense on this understanding of the Textbook categories. Yet many do ask this question because they assume that 'virtue ethics' is not only the name for a particular ethical tradition in Western philosophy but also one of a few basic categories that all ethical theories can be sorted into. I have shown that we might be

better off asking instead about the normative structure, neutrality/relativity, and maximizing status of Confucianism.

In summary, understanding the Textbook categories as family resemblance concepts that mark historical associations may *make sense* of them to some degree. However, it does not *vindicate* the Textbook View because it does not support the claim that these categories track the fundamental divisions in ethics.

5. Concluding remarks

Above I have argued that the Textbook View is a failed paradigm. There is not one fundamental division in ethics giving us three basic options. Instead, there are at least three fundamental divisions giving us a diverse range of theoretical options. If you accept this argument, you might be tempted to conclude that the Textbook categories are illegitimate and best abandoned. Martha Nussbaum takes this line about 'virtue ethics':

> I propose that we do away with the category of 'virtue ethics' in teaching and writing. If we need to have some categories, let us speak of Neo-Humeans and Neo-Aristotelians, of anti-Utilitarians and anti-Kantians—and then, most important, let us get on with the serious work of characterizing the substantive views of each thinker about virtue, reason, desire, and emotion—and deciding what we ourselves want to say. (1999, 201)

Although I am sympathetic to this stance, such a strong claim cannot be justified by the arguments offered here. If the family resemblance picture sketched above is correct then the Textbook categories play a useful role in describing certain historical traditions and spheres of influence. Furthermore, the Textbook categories can be useful shorthand for gesturing towards a particular approach to ethics without precisely marking it boundaries or claiming that it is fundamental. However, both these roles are limited. My arguments against the Textbook View show that even if 'consequentialism', 'deontology', and 'virtue ethics' can do this kind of work, they are unsuited to the more central task of marking the fundamental divisions in ethics. Therefore, whenever we are inclined to use the Textbook terms, we should think carefully about what we are trying to say and ask ourselves whether describing things in terms of the three fundamental divisions articulated here would better convey our point.

One area where this advice may help is in the field of comparative philosophy. The Textbook categories are often applied to moral theories in

non-Western philosophical traditions. For example, in Chinese philosophy some have interpreted Confucianism as a form of virtue ethics and Mohism as a form of consequentialism.[35] Others have disputed these interpretations, arguing that they wrongly impose a Western way of thinking onto traditions with a fundamentally different orientation.[36] Those who take this latter view might be tempted to conclude that the fundamental divisions of Western moral philosophy are inapplicable to non-Western traditions. However, an alternative diagnosis is available. Perhaps the application of 'consequentialism', 'deontology', and 'virtue ethics' to moral theories in Chinese philosophy is questionable because they are historically loaded family resemblance concepts. If so then comparative philosophers ought to move beyond these categories and work with divisions, such as those I propose here, that better capture the fundamental divisions in contemporary analytic normative ethics. This would allow them to better assess whether there are fundamental divisions in ethics with cross-culturally validity. As we saw in §2, this will crucially depend on whether the considerations that motivate the fundamental divisions I propose here are also applicable to assessing the plausibility of moral theories in non-Western traditions. For example, when we evaluate moral theories like Confucianism and Mohism for their overall plausibility, should their ability to capture deontic constraints and special duties, or to satisfy the Compelling Idea be an important part of this evaluation? Focusing on this question seems more fruitful than becoming bogged down in debates about whether calling Mohism a form of 'consequentialism' is illegitimately 'Westernizing' it.

Another benefit of thinking in terms of the three fundamental divisions is that they might illuminate new theoretical possibilities. The Textbook View superficially reduces our options to three. Yet there are many interesting combinations of different stances on the three fundamental divisions that don't fit well with the Textbook categories. Some of these combinations have been neglected and deserve further exploration.

Finally, focusing directly on fundamental divisions without the distracting influence of the Textbook View helps us to clarify what is most important in our moral theorizing. I have presented a strong case for the fundamentality of three divisions given the general concerns of normative ethicists. However, my arguments are not intended to rule out the

---

[35]For example, see Van Norden (2007).
[36]For example, see Ames (2011) and the critical essays in Angle and Slote (2013).

possibility that there are further fundamental divisions which matter in certain localized debates within normative ethics.[37]

Disclosure statement

ORCID

Matthew Hammerton ⬤ http://orcid.org/0000-0003-4649-8206

References

Ames, R. T. 2011. *Confucian Role Ethics: A Vocabulary*. Honolulu: University of Hawai'i Press.

Angle, S., and M. Slote. 2013. *Virtue Ethics and Confucianism*. New York: Routledge.

Annas, J. 1993. *The Morality of Happiness*. New York: Oxford University Press.

Annas, J. 2011. *Intelligent Virtue*. New York: Oxford University Press.

Bentham, J. [1789] 1907. *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press.

Brandt, R. B. 1979. *A Theory of the Good and the Right*. Oxford: Oxford University Press.

Broad, C. D. 1930. *Five Types of Ethical Theory*. New York: Harcourt.

Broome, J. 1991. *Weighing Goods*. New York: Blackwell.

Brown, C. 2011. "Consequentialize This." *Ethics* 121: 749–771.

Burgess-Jackson, K. 2013. "Taking Egoism Seriously." *Ethical Theory and Moral Practice* 16: 529–542.

Chappell, R. Y. 2012. "Fittingness: The Sole Normative Primitive." *The Philosophical Quarterly* 62: 684–704.

Chappell, T. 2014. *Knowing What to Do*. Oxford: Oxford University Press.

Cox, R., and M. Hammerton. 2021. "Setiya on Consequentialism and Constraints." *Utilitas* 33: 474–479.

Crisp, R. 2015. "A Third Method of Ethics?" *Philosophy and Phenomenological Research* 90: 257–273.

Cullity, G. 2015. "Neutral and Relative Value." In *The Oxford Handbook of Value Theory*, edited by I. Hirose and J. Olson, 96–116. Oxford: Oxford University Press.

Cummiskey, D. 1996. *Kantian Consequentialism*. New York: Oxford University Press.

Dancy, J. 1993. *Moral Reasons*. Oxford: Blackwell.

Dancy, J. 2004. *Ethics Without Principles*. Oxford: Oxford University Press.

Dougherty, T. 2013. "Agent-Neutral Deontology." *Philosophical Studies* 163: 527–537.

Dreier, J. 1993. "Structures of Normative Theories." *The Monist* 76: 22–40.

---

[37]For helpful comments on an earlier draft of this article, I would like to thank Steven Burik, Ryan Cox, Andrew Forcehimes, Dan McKay, Thaddeus Metz, Daniel Nolan, Sovan Patra, Tan Sor-Hoon and an anonymous referee from this journal. I am also grateful to audiences at the Australian National University, University of Sydney, Nanyang Technological University, and the 2019 Australasian Association of Philosophy Conference for their helpful input.

Dreier, J. 2011. "In Defense of Consequentializing." *Oxford Studies in Normative Ethics* 1: 97–118.

Dreier, J. 2018. "World-Centered Value." In *Consequentialism, New Directions, New Problems*, edited by C. Seidel, 31–50. New York: Oxford University Press.

Driver, J. 1996. "The Virtues and Human Nature." In *How Should One Live? Essays on the Virtues*, edited by R. Crisp, 111–129. Oxford: Clarendon Press.

Finnis, J. 1980. *Natural Law and Natural Rights*. Oxford: Oxford University Press.

Foot, P. 1985. "Utilitarianism and the Virtues." *Mind; A Quarterly Review of Psychology and Philosophy* 94: 107–123.

Foot, P. 2001. *Natural Goodness*. Oxford: Clarendon Press.

Frankena, W. 1973. *Ethics*. New Jersey: Prentice-Hall.

Garcia, J. 1986. "Evaluator Relativity and the Theory of Value." *Mind; A Quarterly Review of Psychology and Philosophy* 95: 242–245.

Hammerton, M. 2017. "Is Agent-Neutral Deontology Possible?" *Journal of Ethics and Social Philosophy* 12: 319–324.

Hammerton, M. 2019. "Distinguishing Agent-Relativity from Agent-Neutrality." *Australasian Journal of Philosophy* 97: 239–250.

Hammerton, M. 2020. "Relativized Rankings." In *The Oxford Handbook of Consequentialism*, edited by D. Portmore, 46–66. New York: Oxford.

Herman, B. 1993. *The Practice of Moral Judgment*. Cambridge, MA: Harvard University Press.

Hirji, S. 2018. "What's Aristotelian About Neo-Aristotelian Virtue Ethics?" *Philosophy and Phenomenological Research* 98 (3): 671–696.

Hooker, B. 2000. *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*. Oxford: Oxford University Press.

Hurka, T. 2001. *Virtue, Vice, and Value*. New York: Oxford University Press.

Hurka, T. 2003. Moore in the Middle. *Ethics* 113: 599–628.

Hursthouse, R. 1999. *On Virtue Ethics*. Oxford: Oxford University Press.

Jackson, F., and M. Smith. 2006. "Absolutist Moral Theories and Uncertainty." *The Journal of Philosophy* 103: 267–283.

Kagan, S. 2000. "Evaluative Focal Points." In *Morality, Rules, and Consequences*, edited by B. Hooker, E. Mason, and D. Miller, 134–155. Edinburgh: Rowman & Littlefield.

Kamm, F M. 2007. *Intricate Ethics: Rights, Responsibilities, and Permissible Harms*. Oxford: Oxford University Press.

Korsgaard, C. 1996. *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press.

Lazar, S. 2018. "In Dubious Battle: Uncertainty and the Ethics of Killing." *Philosophical Studies* 175: 859–883.

LeBar, M. 2013. *The Value of Living Well*. New York: Oxford University Press.

Louise, J. 2004. "Relativity of Value and the Consequentialist Umbrella." *Philosophical Quarterly* 54: 518–536.

McAleer, S. 2007. "An Aristotelian Account of Virtue Ethics: An Essay in Moral Taxonomy." *Pacific Philosophical Quarterly* 88: 208–225.

McNaughton, D., and P. Rawling. 1991. "Agent-Relativity and the Doing-Happening Distinction." *Philosophical Studies* 63: 167–185.

McNaughton, D., and P. Rawling. 1993. "Deontology and Agency." *The Monist* 76: 81–100.

Mill, J. S. [1861] 1998. *Utilitarianism*. Oxford: Oxford University Press.

Murdoch, I. 1971. *The Sovereignty of Good*. London: Routledge.

Nagel, T. 1970. *The Possibility of Altruism*. Princeton: Princeton University Press.

Nagel, T. 1986. *The View from Nowhere*. New York: Oxford University Press.

Nair, S. 2014. "A Fault Line in Ethical Theory." *Philosophical Perspectives* 28: 173–200.

Nussbaum, M. 1999. "Virtue Ethics: A Misleading Category?" *The Journal of Ethics* 3: 163–201.

Parfit, D. 1984. *Reasons and Persons*. Oxford: Clarendon Press.

Pettit, P. 1997. "The Consequentialist Perspective." In *Three Methods in Ethics*, edited by M. Baron, P. Pettit, and M. Slote, 92–174. Oxford: Blackwell.

Pettit, P., and M. Smith. 2000. "Global Consequentialism." In *Morality, Rules, and Consequences*, edited by B. Hooker, E. Mason, and D. Miller, 121–133. Edinburgh: Rowman & Littlefield.

Portmore, D. 2001. "Can an Act-Consequentialist Theory be Agent Relative?" *American Philosophical Quarterly* 38: 363–377.

Portmore, D. 2005. "Combing Teleological Ethics with Evaluator Relativism: A Promising Result." *Pacific Philosophical Quarterly* 86: 95–113.

Portmore, D. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York: Oxford University Press.

Portmore, D. 2013. "Agent-Relative vs. Agent-Neutral." In *The International Encyclopedia of Ethics*, edited by H. LaFollette, 162–171. Malden, MA: Wiley Blackwell.

Portmore, D. 2016. "Uncertainty, Indeterminacy, and Agent-Centred Constraints." *Australasian Journal of Philosophy* 95: 284–298.

Rawls, J. 1971. *A Theory of Justice*. Cambridge: Harvard University Press.

Ross, W. D. 1930. *The Right and the Good*. Oxford: Clarendon Press.

Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.

Scheffler, S. 1985. "Agent-Centred Restrictions, Rationality, and the Virtues." *Mind; A Quarterly Review of Psychology and Philosophy* 94: 409–419.

Scheffler, S. 1994. *The Rejection of Consequentialism*. Oxford: Clarendon Press.

Schroeder, M. 2007. "Teleology, Agent-Relative Value, and 'Good'." *Ethics* 116: 265–295.

Schroeder, M. 2011. "Ought, Agents, and Actions." *The Philosophical Review* 120: 1–41.

Sen, A. 1982. "Rights and Agency." *Philosophy & Public Affairs* 11: 3–39.

Setiya, K. 2018. "Must Consequentialists Kill?." *The Journal of Philosophy* 115: 92–105.

Sidgwick, H. 1907. *The Methods of Ethics*. Indianapolis: Hackett. seventh edition, 1981.

Slote, M. 1992. *From Morality to Virtue*. New York: Oxford University Press.

Slote, M. 2001. *Morals from Motives*. Oxford: Oxford University Press.

Smith, M. 2003. "Neutral and Relative Value After Moore." *Ethics* 113: 576–598.

Smith, M. 2005. "Meta-Ethics." In *The Oxford Handbook of Contemporary Philosophy*, edited by F. Jackson and M. Smith, 3–30. New York: Oxford University Press.

Smith, M. 2009. "Two Types of Consequentialism." *Philosophical Issues* 19: 257–272.

Vallentyne, P. 1987. "The Teleological/Deontological Distinction." *The Journal of Value Inquiry* 21: 21–32.

Van Norden, B. 2007. *Virtue Ethics and Consequentialism in Early Chinese Philosophy*. Cambridge: Cambridge University Press.