

A Philosophically Informed Virtue Science

In Chapter 1, we laid out the contours of our virtue theory: our general, schematic account of the formal aspects of various virtue traits. We also suggested big picture questions about virtue's importance and development that we think should orient virtue science. In Chapter 3, we say more about the science of virtue and how it builds on, but also offers more than, the pathbreaking VIA research program initiated by Peterson and Seligman (2004). In short, the STRIVE-4 Model does better because it offers an explicitly articulated multicomponent scheme that virtue scientists can use to assess the work that has been done on specific virtues. It also clearly identifies where future work can and should be done. We demonstrate how the STRIVE-4 Model can, in that way, fruitfully guide virtue science by discussing work on the virtue of gratitude and new avenues of research on that virtue that come into view with this approach. In later chapters, we also discuss how it can integrate and advance some existing areas of psychological research, such as moral development and personality psychology.

Before getting into those details, however, we want to highlight the theoretical or philosophic assumptions that lie behind the STRIVE-4 Model and clarify how philosophic work on virtue and specific virtues can best enrich virtue science. Philosophers in a variety of cultures have been debating questions about the nature, value, and development of virtue for centuries, and contemporary philosophic debates on these topics, as with other ones, often involve the development of various incompatible views. Given this dizzying amount of material and the lack of consensus on most topics of contemporary philosophic debate, virtue scientists may be tempted to simply ignore philosophic work on the virtues, or to draw on whatever work they happen to encounter and find inspiring. Although those reactions would be understandable, we think there is a better, more systematic approach that scientists can and should adopt.

Although philosophers seldom reach consensus on questions about virtue and its development (or on anything really!) their collective habits of analysis, argument, and debate enable us to rigorously distinguish and assess the pros and cons of various views. Of course, while many of these topics will interest virtue scientists, not all of them are directly relevant to questions about how to pursue scientific research best and most fruitfully. For example, virtue scientists would probably be interested in philosophic debates about whether a virtuous exemplar would have an abortion or not in some specific situation and in debates about whether what a virtuous exemplar would do is a good guide to what is morally required or not. But those questions about how exemplars act, and the moral relevance of those facts, are issues that scientists can and should set aside when they are thinking about how to fruitfully study virtue and virtue development. Scientists can certainly take an interest in those issues for their own sake and, as we discuss in Chapter 12, the results of virtue science may have a bearing on philosophic debates about how to evaluate and develop competing theories in normative ethics (e.g., Kantian, utilitarian, and virtue ethical), but in this chapter we mainly focus on philosophic debates that can and should inform virtue science. The short answer is that there are three relatively accessible areas of philosophic debate on the virtues that scientists can fruitfully learn about themselves or tap philosophers to learn about.

First, there are numerous claims that philosophers, historical and contemporary, have made about virtue and virtue development. Awareness of these claims will enable scientists to either identify and test new empirically tractable hypotheses or recognize any contentious theoretical background assumptions that they are making. Second, philosophers have long disagreed about how virtue and virtue development are related to human happiness, well-being, or flourishing, both individually and relationally or communally. Awareness of the range of positions and the different conceptions of virtue and well-being in philosophic debates enables scientists to identify and test new hypotheses or to recognize when they are making contentious theoretical background assumptions. Finally, philosophers have offered, and continue to offer, competing conceptual analyses of various virtue concepts and traits; and awareness of those concepts and traits enables scientists to assess and improve scales, measures, and constructs. It also enables the recognition of contentious assumptions that they are making and building into their studies and interpretations of the results.

In the rest of the chapter, we provide some examples and details to illustrate these points. We also highlight the benefits that virtue science

can expect as it becomes an increasingly interdisciplinary project. We start by saying a bit more about the main areas of moral philosophy and some of the interesting philosophic debates, including debates about virtue, that we think scientists can safely set aside. This clarifies our point about many philosophic debates being interesting but irrelevant to scientific inquiry. It also allows us to clarify some disciplinary differences between philosophers and psychologists when it comes to the study of virtue. After that, we turn to the three sorts of philosophic debates that virtue scientists can and should mine either on their own or through interdisciplinary consultation or collaboration.

How Moral Philosophy Transcends the Concerns of Virtue Science

Generally speaking, scientists are interested in studying the kinds of moral traits that show up in the people they study and that human beings in various cultures recognize in ordinary life (not just in scientific contexts) as moral virtues. In conducting studies of these virtue traits, their development, and their impact on human lives and relationships, scientists have to make theoretical assumptions about the relevant virtues. For example, we need to make assumptions about compassion and justice in order to build models and scales, which we use to investigate who is compassionate and just and what promotes or hinders increased compassion and justice. These theoretical assumptions enable scientists to measure virtue traits that ordinary subjects recognize, value, and attribute to themselves and others. As we explain in the next section, however, some philosophic work *can* usefully inform this scientific work because it can help scientists to think creatively and responsibly about the virtue traits they aim to measure, the possible processes of moral trait development, and the forms of well-being or flourishing that virtue might facilitate. But in this section, we want to clarify that many of the main subjects in moral philosophy, and even a good bit of the philosophic work on virtue itself, goes well beyond what is relevant for science.

Much of moral philosophy can be divided into two subjects: normative ethics and metaethics. *Normative ethics* is primarily devoted to questions about moral right and wrong, moral virtue and vice, rationality and irrationality, well-being and suffering, and impersonal good and bad. Normative ethicists develop and argue for different general theories on these broad topics, and they also delve into “applied” ethical debates, such as the debates about the morality of abortion and euthanasia. *Metaethics* in

effect brackets all these normative questions about which things are right, wrong, virtuous, vicious, etc., and turns to questions about the nature and importance of normative properties and the language that is used to talk about them. For example, metaethicists address questions about how morality fits into the natural world, about whether there are objective facts about morality and virtue, and about whether morality or rational moral commitment depends on some sort of supernatural or theistic assumptions. As this makes clear, moral philosophy tackles a variety of important and perennial topics, but it should also be clear that much of moral philosophy has no direct relevance to virtue science. Individual scientists may adhere to various metaethical views, for example about how moral facts fit into the natural world, but it is unlikely that their approach to virtue science hinges on those assumptions, so those assumptions do not necessarily need to (or even can) be empirically assessed. Similarly, virtue scientists are likely to be interested in debates about the morality of abortion and perhaps general questions about moral obligation and how we should respond to wrongdoers. However, moral philosophic work on these topics generally does not bear directly on how we think about virtue traits and study them in a scientific setting, so they can be set aside too.

The main areas of moral philosophy that are directly relevant to virtue science are the parts of normative ethics that focus on virtue and personal good (well-being, happiness, etc.), but, perhaps surprisingly, even some central work on virtue is not relevant to virtue science due to disciplinary differences between philosophy and empirical science. The basic difference here is that while scientists aim to understand, predict, and explain the realistic human moral virtues that are valued in ordinary practice (e.g., by the people they study), philosophers are more interested in developing general theories of ideal virtue. These theories are often revisionary relative to ordinary thought and practice and also have a kind of maximal generality so that they can apply to real or possible beings who have psychologies that differ from human beings (e.g., angels or aliens). Insofar as philosophers argue that we should radically revise ordinary thought or broaden their scope beyond human beings and human psychology, virtue scientists can safely leave their work aside for purposes of scientific inquiry.

As an example of philosophers offering an account of ideal virtue that revises and therefore goes against ordinary thought about moral virtues, consider the doctrine of the unity of the virtues. Some famous moral philosophers such as Plato, Aristotle, and the Stoics argued that we should reject the ordinary assumption that people can have some virtue traits but not others and that sometimes people embody virtues in action but fail to

act in optimal or ideal ways.¹ Against this, they hold that there is a unity of the virtues such that one cannot have one virtue trait and lack any others and that all virtues involve a kind of practical wisdom which ensures that no action that expresses a genuine virtue can be a mistake (all things considered). As proponents of this unity doctrine would admit, these are revisionary views that conflict with ordinary thought and practice, and that are for this reason in need of philosophic defense. In ordinary practice, we often treat virtue traits as modular and fallible, allowing that someone could be compassionate but also a bit cowardly or that someone could act in a kind and humble way that is less than ideal in some circumstances because they are conflict averse. Moreover, if we adopt the strong unity of the virtues view, it is likely that no one you know has *any* virtue traits because, on that view, no one has individual virtue traits such as kindness or courage unless they have an overall virtuous character of the sort that we would normally associate with sages or saints. Of course, there are interesting philosophic arguments that aim to push people to accept these and other revisionary normative views (see Cooper, 1998; Russell, 2009; Wolf, 2007), and scientists may no doubt be interested in the revisionary normative ideals that philosophers argue we should adopt. But they can safely ignore those ideals and arguments insofar as they aim to investigate the common-sense or ordinary kinds of moral traits that we, and the subjects of virtue science experiments, attribute to one another in our day-to-day lives.

In addition to developing revisionary but rationally mandated normative ethical ideals, philosophers also aim to provide maximally general philosophic accounts – roughly maximally general conceptual analyses² – of concepts such as right action, virtue, and well-being. The salient point here is that scientists aim to understand and explain the moral traits of *human beings in this world*, so facts about human evolution and human psychology are relevant to their investigations, and they need not be concerned about what virtues would look like in an alien species with a very different psychology. By extension, they need not be concerned with philosophic debates about the nature or concept of virtue that aim for that kind of larger generality.

To make this concrete, consider the lively debate among philosophers about whether virtue traits must involve positive inner attitudes, such as

¹ For an intelligent and nuanced interdisciplinary discussion of the unity of virtues that suggests a modified and more modest version of this doctrine, see Wright et al. (2021).

² We describe philosophers as providing conceptual analyses because we think it is a helpful term for virtue scientists, but we also note that many philosophers would prefer an alternative description.

emotions and intentions, toward good things such as others' flourishing and preclude positive attitudes toward bad things (e.g., a sadistic desire to cause others pain and suffering). Although ordinary thought and talk about human virtues pretty clearly does assume that virtues must involve good inner states and preclude bad ones, there are philosophers who defend the externalist and broadly consequentialist view that traits are virtues, just in case they tend to have overall good impacts in the local environment and that any kind of inner psychological traits can be a part of a virtue trait (if, that is, it enables positive external impact). To see how this plays out, consider an example given by Driver (2001), the foremost defender of virtue externalism:

Imagine a society that has evolved differently from human society. These creatures, Mutors, have evolved in an extremely harsh environment and have developed unusual strategies for survival. It happens to be the case that for them, beating one's child severely when it is exactly 5.57 years old actually increases the life expectancy of the child by 50 percent ... It is also the case that the only way a Mutor could ever bring himself to so treat a child is to develop an intense pleasure in doing so. So some Mutors have a special trait – they intensely desire to beat children who are exactly 5.57 years old ... On my view this trait would be a virtue. It is an “excellence of character” because it is valuable in that it actually does produce good and a significant social benefit, and the trait is specific enough so as not to produce overwhelming bad consequences. (pp. 55–56)

Of course, Driver realizes that this claim about conceivable but unrealistic cases is bound to strike most of us as counterintuitive, but she thinks this is because our intuitions about what is a virtue or not are trained up in more realistic and actual cases. Our best philosophic account of virtue may differ from the common-sense concept because philosophers aspire to attain kinds of generality and theoretical coherence that are not embodied in common-sense thinking. As she puts it:

[T]o say that human beings are so constructed as to be unable to be virtuous while acting with bad intentions is to *state something contingent*, something about human nature; it is not to state something definitive about virtue unless one can argue for a chauvinistic thesis that moral virtue can obtain only for human beings, and not for any intelligent social creature. A theory of virtue should be broader than this. *It must be conceptually possible to speak of the moral virtues of Mr. Spock.* (p. 56, emphasis added)

In response, defenders of “virtue internalism” agree that our common-sense intuitions about virtue clash with what Driver says about this far-flung case. They also defend general theories that fit and extend the

common-sense assumptions that her theory rejects. Many of them hold, for example, that moral virtues, even if instantiated in imagined Vulcans or Mutors, must involve positive attitudes toward good things and negative attitudes toward bad things. For this reason, even if traits involving sadism happen to lead to good results in some environments, they are nonetheless not moral virtues. Put otherwise, internalists argue that, even in Driver's imagined Mutors case, we should call the traits instrumentally beneficial vices, not moral virtues.

This dispute illustrates the fact that Driver and many other philosophers are interested not just in the concepts of virtue and vice that human beings actually use and that apply to human beings with normal or realistic human psychologies, but in philosophically more general accounts of virtue that apply to even the most artificial and far out cases.³

In addition, philosophers are interested in resolving disputes about which lists and understandings of virtue are best or objectively correct and about how important or valuable virtue is. Although some argue that all rational agents should aim to become as virtuous as possible or at least should develop some specific virtues, others disagree. All of this work is important and interesting, but virtue scientists can safely set it aside while conducting research. As mentioned in Chapter 1, virtue scientists can recognize the importance of respectful internal critique but also be ecumenical and simply recognize and study the diverse lists and understanding of the virtues that they encounter in various cultures and subpopulations. Doing so does not require simply assuming these traits are best, because one can consciously bracket that question and questions about how virtuous people should be. As we discuss more in Chapter 12, the *results* of virtue science can even help us answer questions on those fronts. For example, if the results show that some currently valued trait is conducive to individual or social harm, then that should be taken into account in philosophic debates about keeping that trait on our best list of virtues and debates about developing that trait.

In sum, philosophers aim to develop ambitious and maximally general accounts of virtue and to tackle questions about the best list of the virtues and the importance or value of virtues. These are certainly vital and interesting topics that will interest most virtue scientists, but they can and should be left aside by psychologists and philosophers *while they are pursuing virtue science*. They may safely bracket questions about what the best completely general theory of virtue looks like and focus their attention on

³ See Van Zyl (2018) for an overview of this and related philosophic debates about the virtues.

the realistic virtue concepts that people use in ordinary practice and that apply to human beings with realistic human psychologies.

How Philosophy Can Enrich Virtue Science

In addition to the philosophic work on virtue that is revisionary, ideal, and maximally general, there are areas of philosophic inquiry and speculation in moral philosophy that are relevant to virtue science. As we explain, this philosophic work can enrich scientific inquiry by enabling scientists to either (1) recognize contentious theoretical assumptions that they are making or (2) identify empirically tractable hypotheses that can be scientifically investigated. Naturally, we cannot offer a comprehensive survey of all relevant philosophic work, but we illustrate three main subareas of philosophic research that are relevant. In each case, we give an example, highlight how awareness of various philosophic views can inform virtue science, and suggest that virtue science can be enriched by philosophic material.

Reason, Emotion, and Developmental Pathways to the Moral Virtues

First, there are philosophic debates, both historical and contemporary, about the roles of reason, intuition, and emotion in virtues and their development. There are actually many debates here, in part because “reason” can be used to refer to many different things, but the salient debate here is between philosophers who think that moral virtues such as gratitude and honesty (or at least high-grade versions of these traits) must involve some kind of articulate and discursive evaluative knowledge or deliberation and those who downplay or reject the importance of articulateness and deliberative rationality or wisdom. Philosophers in the second camp do think that virtues involve a kind of intelligent sensitivity to values and the evaluative considerations that are salient in a given situation, but they usually think that this sensitivity is a matter of (native or skilled) intuition or emotional sensitivity, whereas deliberation involves virtue concepts and conceptual articulateness about why responses are virtuous. This deliberation may not be necessary to have virtue traits. For example, Driver (2001) argues that modesty and several other kinds of virtue involve ignorance, not articulate knowledge, and Slote (2010) maintains that moral virtue is centrally a matter of mature empathy, not articulate knowledge or deliberative rationality. We, the authors, favor the former camp of philosophers who hold that discursive, deliberative knowledge is an inherent aspect of at least

many high-grade virtue traits. But we recognize that this is a contentious position and think that virtue scientists who build one view or another into their models and studies should recognize the contentiousness of the assumption that they are making. They should recognize this both because the relevant assumption is a contentious one about the ordinary moral virtue traits that one is trying to study, and because the philosophic debate in this area reflects a diversity of views that one would likely find in different cultures and research participants.

In addition to these debates about the role of articulate knowledge and rational deliberation in virtue traits, there is a historical and contemporary philosophic debate about the processes of moral development – a debate which centers on questions about whether virtue traits are best developed, or can even be effectively promoted, by reflection, habituation, meditation, argument, teaching, or some other means. For example, philosophers in the Aristotelian tradition tend to assume that the pursuit of discursive knowledge and reflection on one's values and life are conducive to virtue development and cultivation (Annas, 2011; Hills, 2015), whereas that position is hotly contested in the Confucian tradition, with some agreeing and others strongly favoring practices of meditation and self-reflection that aim toward acting in accord with one's conscience or innate knowledge of right and wrong (Angle & Tiwald, 2017). And, unsurprisingly, philosophers such as Slote (2010), who give empathy pride of place in their theory of moral virtue, hold that empathy-encouraging processes and not rational reflection or deliberation promote moral virtue most effectively.

As in the previous case, we think virtue scientists may reasonably adopt their own contentious assumptions on these issues when pursuing their work, but that if they do so they should recognize that they are making the assumptions. This recognition of assumptions also means that scientists should be open to others pursuing research on different assumptions.

Even better, we think scientists should try to find ways to empirically test the relevant assumptions where possible. If virtue scientists pursue this third, ambitious option, they can help to settle longstanding and seemingly intractable philosophic debates about virtue and virtue development by converting the core theses of the opposing sides on various debates into empirical hypotheses that can be tested. For example, instead of *assuming* that a virtue trait such as gratitude, honesty, or modesty will be strengthened by reflection on the value of the trait (perhaps by writing about personal experiences or exemplars) or by deliberation about what would count as manifesting the virtue or not, we can treat them as interesting empirical hypotheses to be tested scientifically.

Morgan and colleagues (2017) adopted this approach in their studies of gratitude as a virtue. They empirically examined multiple components of gratitude (knowledge, emotion, attitude, and behavior) and found that each of the components contributed to the explanation of well-being, even when controlling for the other components. Their studies provide an example of testing the value of theoretical concepts (virtues with multiple components) that can be usefully emulated by other scientists.

To see why this is an important point that may not be obvious to all virtue scientists, consider the framework for virtue science recently proposed in Wright and colleagues (2021).⁴ These authors meld an impressive blueprint for empirical investigations, which we discuss and broadly agree with in later chapters, with a broadly Aristotelian theory of virtue. Some parts of the resulting program are ones we endorse because we think that they make only philosophically uncontentious assumptions about realistic, human moral virtues. A prime example is the authors' helpful and plausible distinction between three functional aspects of virtue traits: patterns of input uptake, characteristic social-cognitive processing of inputs, and outputs in the form of "trait-relevant responses." Compassion, for example, might be thought to involve dispositions to notice when others are suffering (input), to feel empathy and a desire to help (social-cognitive processing), and to then act beneficently (output). This three-part model is useful because, like Morgan and colleagues (2017), it encourages conceptual questions about what virtue traits involve, and it can guide empirical studies by highlighting various aspects of a virtue for empirical research to target and measure. Moreover, the authors distinguish various aspects of the three parts and show how they can be assessed in studies in order to measure the dynamic inner structure of virtue traits that are embodied at different times and in different contexts (pp. 34–60, 121–187). The resulting model and the extensions they make to measuring multiple virtues, their interaction, and the ways we can model overall character are insightful and fit in well with our STRIVE-4 approach.

Things become contentious, however, when it comes to Wright and colleagues' claims about practical wisdom and its role in virtue and virtue development. First, the authors take practical wisdom to involve both developed forms of intelligent responsiveness to situations *and* a set of more strongly rational tendencies that involve things such as reflection,

⁴ It is worth noting that this book is a collaboration between two psychologists (Wright, Warren) and a philosopher (Snow). It is an excellent example of the kind of interdisciplinary collaboration we are advocating in this book.

deliberation, and reason-based judgment (pp. 22–30; Snow et al., 2021). In the former vein, they mention (1) situational comprehension or apt perception and (2) sympathy-informed awareness of how it is appropriate to respond to perceived situations and people. In the latter vein, they mention (3) theoretical and practical problem-solving abilities, (4) sound deliberative reasoning about how to achieve chosen ends, (5) wise reflective choice of ends to pursue, and (6) wise reflection on one's values, one's roles and relationships, the kind of person one is and wants to be, and the narrative structure or unity of one's life as a whole. There is nothing inherently problematic about using "practical wisdom" to refer to all six of these abilities or achievements, but to avoid confusion we use "nondeliberative practical wisdom" to refer to items 1 and 2 discussed earlier and "deliberative practical wisdom" to refer to items 3 through 6.

The main point in the current context is that, when Wright and colleagues (2021) suggest that virtue scientists should adopt a framework which *assumes* that deliberative practical wisdom is integral to moral virtue traits, they are in effect suggesting that scientists adopt philosophically contentious assumptions. We think that can be fine, but that scientists need to recognize and flag contentious philosophic assumptions, if and when they make them – ones that some philosophers reject and that we can also expect many ordinary people to reject.

In addition to helping scientists to recognize and flag their contentious assumptions, awareness of relevant philosophic debates enables scientists to sometimes formulate competing empirical hypotheses that can be tested to help settle the underlying philosophic debates. So, for example, instead of assuming that increased theoretical and practical problem-solving abilities correlate with increased gratitude or some other virtue, we think this could be formulated as an interesting hypothesis that could be tested. Philosophers who endorse other approaches to virtue and virtue development (e.g., the hypothesis that increased empathic abilities correlate with increased gratitude) may also suggest complementary hypotheses that are worth testing.

Similar conclusions apply to views on the processes that facilitate moral virtue development. To see this, consider the four functions that Wright et al. take practical wisdom to serve in the manifestation and development of virtue traits: "guiding the action of specific virtues," "regulating interactions among multiple virtues," "regulating virtuous emotions," and "guiding reflection on one's life as a whole" (2021, p. 30).⁵ We believe

⁵ We note the strong similarities in Wright and colleagues' (2021) model of practical wisdom and the Aristotelian Phronesis Model (APM) proposed by Kristjánsson and colleagues (2021). The APM

that Wright and colleagues' Aristotelian model mixes uncontentious and contentious theoretical or philosophic assumptions in ways that can and should be parsed apart. In this case, we believe that the first three factors are uncontentious. They pick out aspects of ordinary moral virtue that are compatible with our model and with a broad array of philosophical positions. In other words, we think that anyone interested in studying good character and promoting virtue can agree that to develop robustly good character traits, many, if not all, developing humans initially need to improve at expressing their traits. This includes accurately hitting the "targets" of the relevant virtues, integrating or regulating their dispositions in cases where different virtues seem to point in different directions, and harboring and expressing virtuous emotions in the most appropriate fashion. For example, when children are still developing a robust trait relating to justice, they may need guidance to express their concern for justice in a productive and compassionate way or they may need to learn from experience in order to develop the courage to act on their sound sense of justice.

Things get more problematic, however, when it comes to the more robustly rational and deliberative aspects of practical wisdom and Wright and colleagues' (2021) apparent assumption that to develop overall virtue, practical wisdom needs to guide "reflection on one's life as a whole."⁶ These Aristotelian claims are philosophically contentious and reflect disagreements that persist at the level of ordinary thought and practice, so we suggest that scientists who are tempted to make them should either flag them as contentious or convert them (and opposing views helpfully expressed in the philosophic literature) into empirical hypotheses.

For an example of how this would play out in detail, consider the authors' useful idea that we could present people with vignettes in which the demands of various virtues are in apparent conflict. Wright and colleagues seem to assume that practical wisdom is the best way for subjects to regulate or adjudicate the apparent interaction or tension in such cases; they say that "practical wisdom must help her navigate the situation" (2021, p. 152) and to assess people's degree of wisdom, and hence virtue, they suggest that experimenters track how people reflectively think about the situations as new nuances are introduced. The assumption seems to be that people who are "able to successfully articulate how and why they

should also be tested in the ways we indicate for the Wright and colleagues' model, and this empirical testing has begun (e.g., Darnell et al., 2022).

⁶ The APM has a similar component, called the "blueprint component," and empirical results were consistent with the blueprint component in a "proof-of-concept" study (Darnell et al., 2022). Further empirical examination of the blueprint component is under way.

would respond to [changes in virtue-relevant stimuli] and why they would respond to them in the way they've reported" (p. 153) are more likely to navigate actual situations of virtue conflict better (cf. Snow et al., 2021, pp. 71–72). But many philosophers would disagree, and assuming one among many possible answers to a conceptual question before we do studies is an unnecessarily contentious approach. That is, we recommend admitting the contentiousness of the assumption at the outset. Of course, we see it as even better to develop and test a set of competing hypotheses.

Similarly, consider Wright and colleagues' (2021) suggestion that overall virtue and virtue development can be best promoted by reflection on one's values, one's roles, virtues one has or aspires to have, and the narrative structure of one's life as a whole (e.g., pp. 27–29, 47–48, 278–291). Once again, we think that a comparison of competing philosophic theories of virtue gives us reason to reduce how much is built into the theory of virtue and to increase our stock of competing empirical hypotheses.

Consider, for example, the lively and complex debate among Chinese neo-Confucian philosophers when it comes to virtue and virtue development (see Angle & Tiwald, 2017, especially chapters 6 and 7). Very roughly, some philosophers such as Zhu Xi (1126–1271) thought virtue could be best developed by the intellectual study of the moral exemplars illustrated in various canonical texts, by comparative rational reflection on one's virtues or vices, and by deliberate attempts to then instantiate the ideals derived from exemplary models. This approach seems to fit hand in glove with the Aristotelian model of virtue and virtue development that Wright and colleagues recommend.

This approach was strongly rejected by philosophers such as Wang Yangming (1472–1529), who thought it was a mistake to divide virtue development into a two-aspect process involving the development of intellectual rational knowledge and then the internalization or application of that knowledge in action. Wang held that human beings have an innate ability to know right from wrong and that the task of virtue cultivation should focus on eliminating things such as selfish desires or self-centered psychological tendencies that block our natural tendency to act on our knowledge of what is morally best. So, he would presumably reject Wright and colleagues' assumption that practical wisdom is needed, or would be useful, for virtue development and suggest an alternative. When it comes to the actual practice of virtue cultivation, Wang would suggest that educators, parents, and others interested in virtue cultivation should adopt a set of strategies that promote awareness and the expression of innate moral knowledge. They should drop any existing attempts to

promote Aristotelian practical wisdom of the sort that involves reflection on exemplars and self-reflection on an ideal that one hopes to instantiate. We think that scientists ought to at least flag their contentious philosophic assumptions and that, ideally, they would work together to identify different hypotheses that are suggested by competing philosophic traditions. This paves the way to differentiate which ones are supported by empirical evidence (keeping in mind that the evidence could be different for different populations or for different moral virtues).

The Moral Virtues and Personal Good

Just as philosophers have long disagreed about the roles of reason and emotion in moral virtue and on the kinds of processes that generate moral development, they have long disagreed about whether moral virtue and its development are conducive to happiness, well-being, flourishing, and meaning. We can broadly classify all these terms under the heading of “personal good.” Philosophers disagree about whether there is any robust connection to be expected between moral virtue and personal good and, if there is a connection, which aspect of personal good is linked to virtue.

We think that the lessons of the previous section on reason and emotion and virtue development apply in similar ways here. First, philosophic work on the different aspects of personal good (happiness, well-being, flourishing, and meaning) can inform and already has been informing work in virtue science and positive psychology. An awareness of philosophic disputes can help scientists to recognize and flag their contentious philosophic assumptions, and also help them to identify new empirical hypotheses to test.

In addition, however, we propose that virtue scientists should engage in interdisciplinary work to better identify various contentious claims that philosophers have made about the connection, or lack of connection, between moral virtue and its development and personal good in its various dimensions. (We suggest many such hypotheses in Chapter 10.) We can illustrate our view here by considering the contentious Aristotelian assumption that wise reflection on the narrative unity of one’s life as a whole, on one’s values, and on the roles and relations one inhabits will be conducive to overall virtue or good character. The theoretical presupposition is that broad evaluative self-reflection will be effective and is perhaps a necessary means to achieving overall good character. This also suggests an assumption that there is a conceptual connection between having overall good moral character and living a good or flourishing life, but both

assumptions are philosophically contentious. Of course, there are various philosophers and psychologists, East and West, who agree with those views, but there are others who defend contrary views, and in this case the bone of contention is not *just* about whether deliberative or reflective rationality will be an effective, or the best, process for the promotion of overall virtue. The larger question is the background assumption that overall good moral character is, in fact, always conducive to living a good and flourishing life or to being a good parent, friend, or human being.

The assumption that moral virtue is conducive to or converges with personal good is certainly common in various longstanding (ancient) philosophic traditions. For example, there are Aristotelians, Hedonists, Stoics, Buddhists, and Confucians who endorse the convergence thesis. But one familiar mark of modern moral philosophy is skepticism about convergence. Western post-Enlightenment thinkers commonly assume that morality and moral virtue are bound to be in tension with personal good and self-interest. In fact, perhaps the central debate among post-Enlightenment philosophers is about whether and why we should be moral *given* our likely forced choice between morality and personal good (Deigh, 2010).

For example, the prime modern anti-moralist, Nietzsche, argued that moral virtue is antithetical to healthy psychological development and agency both in individuals and in larger cultures. Following Nietzsche, we should aim to reject and replace moral virtues and cultural assumptions about their supreme importance or value, thereby freeing human beings from their “false consciousness” about morality (Leiter, 1997). Less extremely, contemporary philosophers such as Williams (1981) and Wolf (1982) suggest that it is rationally permissible to choose our personal good over morality and to refuse to develop saintly forms of moral virtue because doing so would require too much personal sacrifice. Although her work is broadly neo-Aristotelian, Tessman (2005) has argued that people in oppressed groups will often have good reason to develop “burdened virtues” that are morally admirable but costly in terms of personal well-being. Other philosophers, such as Immanuel Kant and his many contemporary followers, would agree with Wolf, Williams, and Tessman that moral virtue often leads to reduced happiness or well-being (at least absent assumptions about a divinely provided afterlife), but they argue that we are rationally bound to act morally and to aim to develop virtuous moral traits.

Finally, a third group of modern moral philosophers assumes that morality and self-interest diverge and express skepticism about our ability to establish what one should favor, rationally speaking. This group

includes Sidgwick, who argues that there is a dualism of practical reason – a conflict between the rational demand to promote our own good and the independent rational demand to act morally – that cannot be resolved (Crisp, 2015). It also includes existentialists such as Sartre, who argues that when moral, prudential, and other evaluative standards conflict, reason cannot tell us what to do. On his view, we are responsible for making a fundamentally nonrational personal choice between morality and personal good when they predictably point us in different directions (Deigh, 2010).

All these views about how to live stand in contrast to the more optimistic views about personal good and moral virtue converging, and they are all live options in contemporary philosophy. So, as in other cases, we think that if psychologists make contentious assumptions about a moral virtue and personal good converging or diverging, then those should be flagged as contentious, and scientists should be glad if others question their assumptions in the interest of scientific open-mindedness and comprehensiveness. Even better, we think that virtue scientists should aim to convert the relevant contentious views into competing empirical hypotheses. Scientists can do this either by studying the philosophic literature themselves or by enlisting philosophers to help them identify contentious assumptions and interesting new empirical hypotheses. That can enable virtue science to be more responsible, ambitious, and exciting.

Conceptual Analysis and the Moral Virtues

The third kind of philosophic research that can helpfully inform virtue science is the increasing range of work on specific moral virtues and what we can think of as the conceptual analysis of moral virtues. Many philosophers are working on historical and contemporary accounts of specific virtues such as modesty, courage, and justice. They tend to offer carefully articulated accounts of what these virtues involve. These philosophers often offer incompatible accounts and argue about which ones are most plausible. One thing these emerging and growing literatures can do is encourage scientists to be aware of possible contentious assumptions they are making about specific virtues when they are building their constructs, scales, or measures.

As an example, consider the literature on humility and modesty. There is a wide variety of views that philosophers have defended but many of them have defined themselves against Driver's (2001) claim that humility and modesty are "virtues of ignorance" which involve people falsely underestimating their excellences or accomplishments, or at least being

ignorant of their excellences or accomplishments. Critics have objected to that account for a variety of reasons (Bommarito, 2018). Proposed alternatives include attention-based accounts on which modest and humble people have true beliefs about their excellences or accomplishments, but on which they are disposed to not draw attention to those excellences or accomplishments, or not disposed to think about them when interacting with others. Regardless of which view readers favor, the general point is that this sort of debate can help scientists to think creatively and critically about the constructs, scales, and studies they might design. They could either incorporate ideas from different philosophic analyses into their studies, or they could posit different aspects and kinds of humility and modesty and conduct studies to understand their relations or differences. Another obvious option here would be to explore the tacit assumptions about virtue concepts that various populations make and that might be affecting things, such as self-report.

Summary

In general, we have in mind the fact that whereas philosophers prize conceptual clarification and exquisitely precise statement and argument, scientists prize constructs, scales, and measures that can be validated and lead to actual research results. Scientists must also design things in light of the fact that their research participants cannot be expected to attend to fine-grained conceptual nuances. As these points suggest, there are good methodological reasons for scientists to adopt theoretical assumptions that are bound to seem somewhat unclear and ham-handed to philosophers and for philosophers to split hairs in ways that scientists are bound to find useless. So, when philosophers aim to help scientists build better models, constructs, and measures, or scientists look to philosophy for greater clarity, they should take care to remember that different disciplinary aims can make this kind of interdisciplinary work frequently trying and inevitably somewhat difficult. In the end, however, we have been arguing that virtue scientists can indeed gain a lot from an acquaintance with philosophic work, both by recognizing contentious assumptions that they may be making and by broadening the scope of their empirical hypotheses to cover a range of interesting issues that have long been left to philosophic speculation.

In this chapter we have addressed questions about philosophy's relevance to virtue science. We have clarified that much of moral philosophy (normative ethics and metaethics) is not relevant to virtue science, even

though it may interest scientists. We have also drawn attention to some aspects of the philosophic debate that can usefully inform virtue science. Having clarified how philosophy can improve virtue science, it is natural to wonder about whether and how virtue science can enrich philosophic debates. After we have further developed our unifying STRIVE-4 framework in Parts II and III, we return to the debates about if and how virtue science can inform or should constrain moral philosophy or not (Chapter 12). For now, we hope to have clarified the interdisciplinary resources that have helped to inform our STRIVE-4 Model and that we think other virtue scientists can fruitfully mine in their own research.