

Why we should both accept and reject hybrid views

Abstract

In this paper I argue that a two-level moral view, with a monist view at the theoretical level and a hybrid view at the practical level, is an attractive alternative to one-level monist and hybrid views. For example, both utilitarianism and rights theory, on a particular interpretation, imply a moderate “Kantianism for people, utilitarianism for animals” in practice. This kind of view preserves the benefits of monist views, since it allows for simplicity, power, and comparability at the theoretical level of morality. It also preserves the benefits of hybrid views, since it honors the concrete, contextual nature of our lived moral experience at the practical level of morality. I also argue that this kind of two-level view is, to its credit, much more “pro-animal” in practice than the traditional “Kantianism for people, utilitarianism for animals.”

1. Introduction

Many people accept hybrid views in ethics. That is, many people accept that different, seemingly incomparable moral theories can be true in different contexts. For example, consider a hybrid view suggested by Robert Nozick in *Anarchy, State, and Utopia* (1974): Kantianism for people, utilitarianism for animals. On this view, we should take a Kantian approach to our treatment of people (i.e., moral and political agents), by treating people as ends in themselves. And, we should take a utilitarian approach to our treatment of animals (i.e., moral and political patients), by maximizing aggregate well-being for

animals. On this view, for instance, it would be wrong to kill one person to save five people (or five animals). But it might be permissible to kill one animal to save five animals (or five people).

Hybrid views have one clear benefit: They accommodate our intuitions that different moral theories make sense in different contexts. However, hybrid views also have several costs. They lack theoretical virtues such as simplicity, power, and comparability. They also, relatedly, have intuitively implausible implications in many cases. For example, in cases where treating people as ends in themselves conflicts with maximizing aggregate well-being for animals, we think that there is a fact of the matter about what we should do all things considered. How can a hybrid view such as “Kantianism for people, utilitarianism for animals” explain this idea? In light of these and other costs, it is worth asking if we can find a moral theory that preserves the benefits of hybrid views while mitigating the costs.

Fortunately, I believe that we can find a moral theory that has these features. If we make a distinction between the theoretical and practical levels of morality, then we will see that we can accept a monist view at the theoretical level and a hybrid view, such as a moderate “Kantianism for people, utilitarianism for animals,” at the practical level. This approach preserves the benefits of hybrid views while mitigating the costs. At the theoretical level, it allows for simplicity, unity, and comparability. At the practical level, it allows for complex, contextual moral decision making. Moreover, I believe, this kind of moral theory is more than a mere possibility. It follows naturally from monist theories such as utilitarianism and rights theory (on a particular interpretation), once we consider what it takes to apply them in practice.

In this paper, I will motivate the idea of accepting a monist view at the theoretical level of morality and a hybrid view at the practical level of morality, focusing on a moderate “Kantianism for people, utilitarianism for animals” as an example. I begin (in §2-3) by discussing the difference between monist and hybrid views, as well as the difference between theory and practice. I then (in §4-5) show why both utilitarianism and rights theory (on a particular interpretation) can imply hybrid views in practice. I then (in §6) examine differences between humans and nonhumans that might affect how we should treat them in practice. Finally (in §7), I show why utilitarianism and rights theory, coupled with these differences, might imply a moderate “Kantianism for people, utilitarianism for animals,” with many caveats.

2. Monist and hybrid views

Monist and hybrid views exist on a spectrum of abstractness and generality. At one end of the spectrum, some theories identify a single, general value, rule, or virtue from which all other values, rules, and virtues derive. At the other end of the spectrum, some theories identify an indefinite number of values, rules, and virtues that apply in particular situations. In between these extremes, there are theories with varying degrees of abstractness and generality. I want to start by considering some representative examples of each kind of theory (without attempting to cover the full spectrum of possibilities), along with their pros and cons. This will set up my suggestion, in later sections, that that a two-level view, with a monist and a hybrid level, can accommodate our moral judgments better than fully monist or hybrid views can.

Classical utilitarianism, Kantianism, and Aristotelian virtue theory are all, as I interpret them, classic examples of monist views. They all involve a simple, general moral aim from which all other moral aims derive. For utilitarianism, this aim is to maximize utility.¹ For Kantianism, this aim is to act only on maxims that you can will as universal laws.² For Aristotelianism, this aim is to flourish as the kind of being that you are.³ In each case, we might find that many goods, duties, and virtues derive from this simple, general foundation, since we might need to pursue many goals, follow many rules, and cultivate many habits in order to accomplish these aims in practice. (More on this point in the next section.) But ultimately, all these further goods, duties, and virtues have only extrinsic, or instrumental, moral value.

Other views can be more complex, in at least two ways. Some views can remain fully consequentialist, deontological, or virtue theoretic, while involving multiple basic goods, duties, or virtues. For example, some people hold that there can be multiple basic goods, such as pleasure, liberty, and equality.⁴ Some people hold that there can be multiple basic duties, such as duties prohibiting lying, stealing, and killing.⁵ Some people hold that there can be multiple basic virtues, such as honesty, respect, and compassion.⁶ These views can vary in complexity, depending on how many basic goods, duties, or virtues they involve. For example, a consequentialism involving two basic goods is simpler than a consequentialism involving four, but more complex than a consequentialism involving one.

¹ Bentham 2018.

² Kant 2012.

³ Aristotle 2016.

⁴ For an example of someone who accepts that there can be multiple basic goods, see Chang 1997.

⁵ For an example of someone who accepts that there can be multiple basic duties, see Ross 1998.

⁶ For an example of someone who accepts that there can be multiple basic virtues, see Swanton 2005.

Other views – which I am calling hybrid views – can combine consequentialist, deontological, and/or virtue theoretic. For example, when Nozick suggests that Kantianism might be the basis of our duties to people and that utilitarianism might be the basis of our duties to animals, he is suggesting that a hybrid view is true.⁷ When Scanlon suggests that contractualism might be the basis of our duties to people and another theory might be the basis of our duties to animals, he is doing the same.⁸ As before, hybrid views can vary in complexity, depending on how many theories they combine, and how many goods, duties, and virtues each theory involves. “Kantianism for people, utilitarianism for animals” is simpler than many hybrid views, but more complex than either Kantianism or utilitarianism for everyone.

The pros and cons of each kind of theory are clear. On one hand, monist views are simple and powerful. They also allow for comparability in morality, since they provide us with a single basic moral value, duty, or virtue that can help us to resolve conflicts between derivative moral values, duties, and virtues. However, monist views seem too abstract and general to honor the concrete, contextual nature of our lived moral experience. They also, relatedly, seem implausible in many cases. For instance, many people have the intuition that multiple goods, duties, and virtues are all intrinsically valuable, as well as that different goods, duties, and virtues apply in different contexts, such as in the case of “Kantianism for people, utilitarianism for animals.” These intuitions seem to militate against monist views.

On the other hand, hybrid views honor the concrete, contextual nature of our lived moral experience. They also vindicate the plausible idea that there are multiple basic

⁷ Nozick 1974.

⁸ Scanlon 1998, 179

goods, duties, and virtues, all of which apply in different contexts. However, hybrid views lack simplicity, power, and comparability. They provide us with no explanation of how these basic goods, rules, and virtues relate to each other, or of what we should do when they conflict. As a result, hybrid views also seem implausible in many cases. For example, suppose that we have to choose between saving a person and saving 10,000 dogs. Intuitively, there is a fact of the matter about what we should do all things considered. But without a “covering value” that allows for comparability, “Kantianism for people, utilitarianism for animals” is silent on this question.

This framing of the pros and cons of monist and hybrid views is standard but, I think, potentially misleading. When we frame the debate this way, we think that we can accept only one kind of moral theory, and so it all comes down to what features of a theory matter most to us. To what degree do we want a theory that can provide simple, powerful moral explanations and resolve conflicts between values? And to what degree do we want moral theories that honor the concrete, contextual nature of our lived moral experience and vindicate the idea that there are different basic values that apply in different contexts? The more we prefer the former virtues, the more we will tend toward simple monist views. And the more we prefer the latter, the more we will tend towards complex hybrid views.

I think that this framing is potentially misleading because it conflates theory and practice. For example, it suggests that if we accept a particular value in theory, then we must always use this value to make decisions in practice. Conversely, it suggests that if we use different values to make decisions in practice, then we must accept a hybrid view in theory. But both of these suggestions can be false. We can accept a relatively simple

monist view in theory, where simplicity and power are more important. And we can accept a relatively hybrid view in practice, where intuitive plausibility is more important. Indeed, I will argue – using “Kantianism for people, utilitarianism for animals” as an example – that such combinations are natural. As a result, we can turn an *either-or* situation into a *both-and* situation.

3. Two-level theories

Many ethicists make a distinction between the theoretical and practical level of morality. We have different names for this distinction – the theoretical and practical levels, the critical and intuitive levels,⁹ decision procedures and criteria of rightness,¹⁰ and so on – but in each case, the distinction is basically the same. At the theoretical level of morality, we ask what makes right actions right. At the practical level of morality, we ask how we can reliably perform right actions in everyday life. This distinction is important because we might need to accept different moral principles at each level. Simply put, applying a principle is not always the best way to conform to that principle. If you want to be happy, then you might sometimes need to pursue goals other than your own happiness. The same might be true for morality as well.

The reason why we need to make a distinction between theory and practice is that humans are not cognitively or motivationally able to consistently apply abstract, general moral principles in everyday life. One limitation concerns the complexity of rules. Abstract and general rules might be simple in theory (since they can explain a wide range

⁹ Hare 1981.

¹⁰ Brink 1989: 256.

of phenomena), but they are complex in practice (since we need to do extra work to determine what follows from these rules in particular cases). Moreover, insofar as abstract and general rules are open to interpretation, there is a risk that bias will distort our application of these rules in practice. For these and other reasons, if we want to reliably comply with an abstract and general rule in practice, then we might need to follow a more concrete and contextual set of rules instead.

Another limitation concerns the role of character and environment in human behavior. We might make *some* decisions explicitly, by asking what we have most reason to do. But we make many *other* decisions implicitly, by doing what seems natural or responding to social or environmental cues. Thus, if we want to reliably act ethically, then we need to consider more than how to act well when we do, in fact, deliberate. We also need to consider how to act well when we do not. For example, we need to cultivate virtuous characters so that we can act well insofar as our character is responsible for our behavior. We need to cultivate good communities so that we can act well insofar as we respond to social cues. And we need to construct good environments so that we can act well insofar as we respond to environmental cues.

The upshot of these (and other) considerations is that *even if* we accept a relatively monist view in theory, we should *still* accept a relatively hybrid view in practice. The more abstract and general a moral principle is, the less useful this moral principle will be for creatures like us in everyday life. For better or worse, we need to think about what to do in terms of concrete, contextual goods and rules, as well as control our actions in other, more indirect ways, such as by creating social and physical environments that bring out the best in us. As a result, in order to reliably promote any

particular good, follow any particular rule, or cultivate any particular virtue, we will likely need to promote a wide range of goods, follow a wide range of rules, and cultivate a wide range of virtues in different contexts in practice.

Of course, these ideas are familiar to indirect, or sophisticated, consequentialists, who have long held that ethics has this structure.¹¹ They think that we should accept the principle of utility in theory, and that we should follow whatever roles, values, rules, habits, and so on will, in fact, maximize good outcomes in practice. Moreover, many consequentialists reject the idea that the correct moral theory must be publicly acceptable. They think that the point of a moral theory is to be correct, not to be accepted or applied in everyday life.¹² So, for instance, if we could maximize good outcomes by burning every book that Bentham, Mill, and Sidgwick ever wrote and telling everyone (including ourselves) that Kantianism, virtue theory, care theory, or the Ten Commandments is true instead, then we morally ought to do that.

However, it is important to note that a two-level morality is available to non-consequentialists too. For instance, in the same way that it is not always possible to maximize good outcomes by consciously and deliberately attempting to do so, it is not always possible to act on universal laws or flourish as the kind of being that we are by consciously and deliberately attempting to do so. Granted, a non-consequentialist might believe that the correct moral theory must be publicly acceptable. But this is consistent with two-level morality. After all, we can act with a background knowledge of the correct moral theory even if we are not thinking in terms of this theory in everyday life, in the

¹¹ See, for example, Sidgwick 2014, Railton 1984, and John and Sebo forthcoming.

¹² Parfit 1984: 24.

same way that we can act with a background knowledge of string theory even if we are not thinking in terms of string theory in everyday life.

When we make a distinction between the theoretical and practical levels of morality, we see that there is not necessarily a need to choose between monist and hybrid views. At the theoretical level, we should expect theories to be relatively abstract and general, since the point of a theory at this level is to provide us with simple, powerful moral explanations and prescriptions. So, it can be fine for a theory to seem implausible at this level, since we should have no expectation that the first principles of morality will be part of ordinary moral reasoning, or that they will reflect our lived moral experience. Thus, we should at the very least think that virtues such as simplicity, power, and comparability are more important than virtues such as intuitive plausibility at the theoretical level of morality.

On the other hand, at the practical level, we should expect theories to be relatively concrete and contextual, since the point of a theory at this level is to empower us to reliably make correct moral decisions. And, it can be fine for a theory to involve a wide range of incomparable roles, values, duties, and virtues at this level, since we have no reason to expect that the complex set of norms that we should follow in everyday life will always be unified in practice, or will always reveal what to do when they conflict in practice (since the practical costs of this extra information might sometimes outweigh the practical benefits). Thus, we should at the very least think that virtues such as intuitive plausibility are more important than virtues such as simplicity, power, and comparability at the practical level of morality.

4. Utilitarianism

Consequentialist criteria of rightness, together with facts about the world, can imply decision procedures that involve consequentialist as well as non-consequentialist norms. To illustrate this point, I will focus on utilitarianism, which holds that we morally ought to maximize well-being in the world. To be clear, I think that this analysis extends to other kinds of consequentialism too, since the features of utilitarianism that call for a partly non-consequentialist decision procedure are present in other kinds of consequentialism as well. But since my main goal in this section is simply to motivate the idea that a consequentialist criterion of rightness *can* imply a partly non-consequentialist decision procedure, I will focus on utilitarianism as a case study and leave it open how far this analysis extends.

In theory, utilitarianism is highly demanding. Every time we act, we have a moral duty to perform the action that maximizes aggregate value for all sentient beings from now until the end of time. This means that there is no room in utilitarianism for “agent-centered prerogatives”.¹³ You might want to devote your career to a particular cause, as well as to create space in life for particular projects and relationships. But utilitarianism implies that you are morally permitted to do so if and only if you can do the most good possible as a result. Granted, there might be *some* cases where you can do the most good possible by doing what you love for a living or by investing in your own projects and relationships. But these cases might be rare. In short, utilitarianism sets a nearly impossibly high standard for right action.

¹³ Scheffler 1982

Additionally, in theory, utilitarianism is highly unrestrictive. We are morally required to maximize aggregate well-being by any means necessary. This means that there is no room in utilitarianism for “side constraints.”¹⁴ For example, if you need to harm one to help five, then you should. If you need to kill one to save five, then you should. And so on. Granted, there might be many cases where harming and killing people does not, in fact, do the most good possible. For example, even if harming one to help five has good results in the short term, it might have bad results in the long run. Still, there is no guarantee that the results of these actions will *always* be bad in the long run. And if and when the results are good in the long run, these actions are not only morally permitted but morally required according to utilitarianism.

However, utilitarianism works differently in practice than in theory, owing to our cognitive and motivational limitations. As we have seen, we lack the information and rationality necessary to estimate the impacts of every possible course of action on every possible sentient being. Many of us also lack the motivation necessary to regularly sacrifice our own projects and relationships for the sake of the greater good. As a result, if we want to do the most good possible in practice, then we should not attempt to maximize utility all the time. Instead, we should reason in the “cool hour of deliberation” about which simple, concrete, and contextual rules, if generally followed, will allow us to promote good outcomes sustainably. And then we should focus for the most part on following those rules in everyday life.

Moreover, as we have seen, while we perform many actions as a result of deliberating about what to do, we also perform many actions as a result of habit, social

¹⁴ Nozick 1974

cues, or environmental cues. Insofar as we do, we need to spend time and energy cultivating habits, building communities, and building environments that will bring out the best in us. Additionally, insofar as we share these structures with others, we need to think not only about what will bring out the best in us but also about what will bring out the best in everyone. This requires individual and shared improvisation, experimentation, and coordination. That is, it requires us to collectively move toward a relatively simple, concrete, contextual, sustainable set of shared norms that we can all internalize, and to individually deviate from these rules only when the stakes are high.

As a result of these cognitive and motivational considerations, utilitarians should incorporate at least some non-consequentialist values into their decision procedure. First, utilitarians should regard morality as less demanding in practice than in theory. Insofar as we need to work on preferred causes or invest in personal projects and relationships in order to make our work sustainable, we morally ought to do so. Moreover, insofar as we decide what to do implicitly, we morally ought to condition ourselves to have character traits that naturally lead to good decisions. This means not only that we should accept agent-centered prerogatives to a degree, but also that we should accept care theory, which emphasizes investing in caring relationships, and virtue theory, which emphasizes cultivating virtuous characters, to a degree.

Second, utilitarians should regard morality as more restrictive in practice than in theory. We need to make sure that we follow good rules and cultivate good characters, so that we can naturally perform good actions even when we lack the ability to think critically about what to do or the motivation to act on these critical judgments. Relatedly, we should be generally skeptical about cost-benefit analyses that purport to show that we

should, say, kill one to save five in a particular situation. This means that we should accept a presumption in favor of side constraints in practice. This presumption is necessary to internalize the rules and habits that maximize utility in the long run. It is also necessary to avoid tragic moral mistakes that can easily result from our ignorance, irrationality, and natural bias in everyday life.

For many utilitarians, the result is a consequentialist criterion of rightness and a partly consequentialist, partly non-consequentialist decision procedure. We morally ought to maximize utility in theory. And, in order to accomplish that aim in practice, we should reason in the “cool hour of deliberation” about what roles, rules, and habits will allow us to do the most good possible in particular contexts. We should then mostly work to play those roles, follow those rules, and cultivate those habits in those contexts. Granted, if we have very good reason to believe that violating these norms will do much more good than harm, then we might be required to violate them for the greater good in practice. But otherwise we should follow them, as an investment in a way of life that will do the most good possible in the long run.

5. Rights theory

Similarly, I believe that non-consequentialist criteria of rightness, together with facts about the world, can imply decision procedures that involve consequentialist as well as non-consequentialist norms. To illustrate this point, I will focus on a version of rights theory that holds that we all have a right to set our own ends in life and a duty to allow

others to do the same.¹⁵ On this view, we are not morally permitted to harm the few to help the many, though there can be exceptions to this rule that we will discuss in a moment. And if we do harm others, then we are generally required to attempt to reduce or repair these harms. As before, I believe that my analysis here can extend to other kinds of non-consequentialism too, but I will focus on this kind of non-consequentialism as a case study and leave it open how far this analysis extends.

In theory, rights theory is not as demanding as utilitarianism. We are not morally required to pursue a particular goal (much less a maximally impartially benevolent goal). Instead, we are morally permitted to pursue our own goals in life, provided that we allow others to do the same. This means that there is room in non-consequentialism for agent-centered prerogatives. If you want to devote your career to a particular cause, then you are morally permitted to do that, provided that the cause is good. Similarly, if you want to create space in life for personal projects and relationships, then you are morally permitted to do that as well, provided that you treat everyone involved with respect. Unlike in the case of utilitarianism, there is no requirement that you attempt to select the roles that allow you to do the most good possible.

Additionally, in theory, rights theory is restrictive. It holds that some actions are morally wrong even if you can do more good than harm by performing them. For example, if people have a right not to be killed, then that means that you are not morally permitted to kill someone even if you can save five people as a result. This means that there is room in non-consequentialism for side constraints. Granted, some kinds of rights theory establish “harm thresholds” above which such sacrifices are “necessary” and, so,

¹⁵ This interpretation of rights theory partly draws from the interpretation of “restricted deontology” that Shelly Kagan develops in 2019.

permissible. On these theories, while you might not be permitted to kill one to save five, you might be permitted to kill one to save, say, one hundred, one thousand, or one million. Either way, in ordinary circumstances we have a moral duty to pursue our goals only in ways that respect the rights of all involved.

However, as with utilitarianism, rights theory works differently in practice than in theory, owing not only to our cognitive and motivational limitations but also to the scale of harm in the world and our relationships with harm in the world. First, consider our relationships with harm in the world. Many rights theorists hold that if we are complicit in unnecessary harm, then we have a moral duty to reduce or repair this harm. And of course, we are all complicit in a lot of unnecessary harm. We are all deeply entangled in harmful and oppressive social, political, and economic systems, and, as a result, many of us contribute to, benefit from, and have the power to address harm and oppression for humans and nonhumans alike. Thus, many of us have a non-consequentialist duty to attempt to reduce and repair this harm and oppression, at least to a degree.

Now consider the scale of harm in the world. As we have seen, some rights theorists hold that there can be “harm thresholds” above which we are permitted, if not required, to harm the few for the sake of the many. Again, different rights theorists might set different thresholds: For some, it might be that we should harm or kill one in order to help or save one hundred. For others, it might be a thousand, a million, or a billion. Either way, we now live in a world where our choice situations might involve such stakes. Granted, we might not regularly be in a position to help or harm millions *individually*. But we are regularly in a position to help or harm millions *collectively*, and to be

complicit in these actions individually. In these cases, we might sometimes need to sacrifice the few for the sake of the (very) many after all.

As a result of these facts about the scale of harm and our relationship with harm in the world, rights theorists should incorporate at least some consequentialist values into their decision procedure. First, rights theorists should regard morality as more demanding in practice than in theory. Given how many harmful systems everyone is complicit in, we might have to think at least partly in consequentialist terms in order to determine how to most effectively reduce and repair this harm. For example, we should consider the scale, neglectedness, and tractability of the problems that we contribute to, benefit from, and have the power to do something about, and we should select the problems with respect to which we can do the most good possible. Granted, our reasons for doing this might be different, but the end result might be similar.

Second, rights theorists should regard morality as less restrictive in practice than in theory. If we accept a harm threshold at all, then we might find that this threshold is regularly exceeded in practice. For example, many nations now have populations with hundreds of millions of humans and hundreds of trillions of nonhumans. Additionally, state actions impact orders of magnitude more humans and nonhumans in other nations and future generations. And, many individuals participate in these actions (as well as in other, comparable collective actions). There is no way to decide what to do at this scale without making trade-offs. And while, for instance, harming one to help five (or one million to help five million) might not be acceptable in these cases, harming, say, one to help a million (or one million to help a billion) might be.

For many rights theorists, the result is the same as for many utilitarians: a non-consequentialist criterion of rightness and a partly consequentialist, partly non-consequentialist decision procedure. We are morally permitted to live and let live in theory. And, in order to accomplish that aim in practice, we should reason in the “cool hour of deliberation” about what roles, rules, and habits will allow us to reduce or repair our complicity in harm most effectively in particular contexts. We should then mostly work to play those roles, follow those rules, and cultivate those habits in those contexts. Granted, if we have very good reason to believe that violating these norms will do much more good than harm, then we might be required to violate them in practice. But otherwise we should allow rights to trump results.

6. Humans and other animals

I now want to motivate the following idea. Utilitarians and rights theorists (on the interpretation described above) should both endorse a decision procedure that involves relatively Kantian duties to humans, relatively utilitarian duties to insects, and various balances for animals in between these extremes. This will provide indirect, though imperfect, support for a moderate “Kantianism for people, utilitarianism for animals” in practice (though this hybrid decision procedure will be much more “pro-animal” than the traditional Nozickian hybrid view). I will start, in this section, by discussing some of the facts about human and nonhuman animals that will be relevant here. I will then, in the next section, argue that both utilitarianism and rights theory, combined with these facts, imply this hybrid decision procedure.

There are many kinds of facts concerning humans and other animals that are relevant here. Some concern what humans and nonhumans are like. Others concern how many humans and nonhumans there are, and how diverse humans and nonhumans are. Still others concern contextual and relational matters: We are capable of systematically harming and helping other animals, but we are capable of harming and helping different animals in different ways. As we will see, the upshot is that we should think of our duties to some animals in more Kantian terms, our duties to other animals in more utilitarian terms, and our duties to other animals in more balanced terms. Moreover, while we can aspire to more of a Kantian “Kingdom of Ends” in a mostly human community, we must aspire to a more of utilitarian ideal in a multi-species community.

1. Welfare. Some animals can have a higher capacity for welfare than others. For a clear example, consider the difference between humans and ants. Individual humans have a higher capacity for welfare than individual ants *synchronically*, since we have more complex brains and nervous systems, and so we can experience more happiness or suffering at any given time. Individual humans have a higher capacity for welfare than individual ants *diachronically*, since we have longer lifespans, and therefore we can experience more happiness or suffering over the course of our lives. And of course, this is not a binary difference between humans and other animals. Rather, it is a spectrum of difference, with humans near (though not necessarily at) one end of the spectrum and ants near (though not necessarily at) the other end.

This difference can affect the strength of our individual interests, needs, and rights. If a particular human has a higher capacity for suffering than a particular ant, then we might think that the human has a stronger interest in avoiding suffering, as well as a

stronger right to avoid suffering. As a result, insofar as we are reasoning as consequentialists, we might think that we should prioritize the human over the ant all else equal, since the human has more at stake, and thus prioritizing the human contributes more to aggregate utility all else equal. And, insofar as we are reasoning as rights theorists (who accept harm thresholds), we might think that we should prioritize the human over the ant all else equal, since the human has stronger welfare rights, and thus the bar for violating their welfare rights is higher.

2. *Agency*. Some animals can also have a higher capacity for agency than others. While many animals are capable of setting and pursuing goals, only humans are capable of setting and pursuing goals by thinking about what we have reason to believe, desire, and do. As a result, only humans can think about how we rationally, morally, and politically ought to treat ourselves and others. Additionally, while many animals have an interest in goods such as education (since, for instance, many animals depend on learning to set and pursue goals effectively, particularly in novel contexts), humans generally have a stronger interest in these goods (since humans generally depend on these goods to set and pursue goals effectively more than other animals do), as well as an interest in additional goods, such as freedom of thought and speech.

This difference can affect the scope and strength of our rights and responsibilities. First, many humans are rational, moral, and political agents. We have a right, and a responsibility, to participate in creating a shared set of laws that can allow everyone to live and let live peacefully. We also have a right, and a responsibility, to set and pursue goals within the confines set by these laws. Second, insofar as humans have more, and stronger, interests related to agency than other animals, we will have more, and stronger,

rights related to agency as well. For example, insofar as humans have a stronger interest in education than, say, mice, we will have a stronger right to education. And, insofar as we have an interest in goods such as freedom of thought and speech, we might have additional rights to such goods.

3. *Numbers*. There are about eight billion humans in the world. Meanwhile, there are *quintillions* of nonhumans in the world. While demographics vary regionally, we can expect that any particular region (for instance, any particular city, state, or country) contains orders of magnitude more nonhumans than humans. Moreover, with some species, such as many K-selected species (species whose members have large bodies, long life spans, and low reproduction rates), we might be able to treat members of these species as individuals at scale. But with other species, such as many r-selected species (species whose members have small bodies, short life spans, and high reproduction rates), we might be less able to do this at scale, and so we might have no choice but to interact with most members of these species in abstract, general terms.

This difference affects what kind of moral and political community we can build. The larger our community is, the less we can interact with everyone as an individual, and the more we have to interact with community members in the aggregate. Likewise, the higher the ratio between nonhumans and humans in our community, the more humans, as moral and political agents, will have to sacrifice to build a just society. This means that we have two options. The first is that we can engage with everyone in relatively impersonal terms. The second is that we can engage with some individuals in more personal terms (such as humans and some nonhumans) and others (such as other

nonhumans) in more impersonal terms. While the first approach might be best in principle, the second approach might be strategically necessary in practice.

4. *Diversity*. There is a lot of diversity within species. But there is much more diversity across them. And while diversity brings many benefits, it also brings many conflicting needs. For example, no matter what climate we have, this climate will benefit some animals and harm others. And, no matter what infrastructure we build, this infrastructure will benefit some animals and harm others. Additionally, many animals survive by harming and killing other animals. This can be true within species; for example, some animals eat their mates or offspring. It can also be true across species. Predator species survive by eating other animals. Parasitic species survive by living on, or in, other animals. And of course, many animals harm or kill other animals for other reasons as well, such as self-defense, other-defense, or play.

This difference again affects what kind of moral and political community we can build. Basically, the more conflict we have, the less able we are to build a “Kingdom of Ends” in which everyone can live and let live with the kind of liberty and security to which humans are accustomed. This means that we once again have two options. The first is that we can accept that our community will involve massive amounts of conflict and attempt to manage this conflict as well as possible. The second is that we can construct a relative Kingdom of Ends for some individuals (such as humans and some nonhumans) and then attempt to manage conflict as well as possible among other individuals (such as other nonhumans). Once again, the first approach might be best in principle, but the second might be strategically necessary in practice.

5. *Tractability*. Finally, and relatedly, we are generally able to support individual humans more effectively than individual nonhumans. We have already considered some reasons why; for instance, there are fewer members of any particular species than across them and less conflict within any particular species than across them. But there are other reasons as well. For instance, we know more about how to support humans than nonhumans, and we have a better infrastructure in place for supporting humans than nonhumans. This is partly our fault: We have prioritized our own needs throughout our history as a species. But it is partly not our fault: We can more easily learn about beings like us than about beings unlike us, and so our uncertainty about human welfare, while still high, is lower than our uncertainty about nonhuman welfare.

This difference once again affects what kind of moral and political community we can build. When we have not only more need and conflict but less tractable need and conflict in our community, our prospects for constructing the kind of Kingdom of Ends to which people are accustomed shrink even further. We once again either need to accept that our moral and political community must involve much more need and conflict than we had hoped, or need to create a moral and political community with less need and conflict for some and more for others. And while we *might* be able to call either of these outcomes a Kingdom of Ends (depending on how we interpret the idea), we will have to accept that such a society looks very different in practice than Kantians imagined that a Kingdom of Ends might look in theory.

6. *The Sixth Extinction*. The harms that humans cause to other animals are well-documented but worth emphasizing. We breed, raise, and kill more than 100 billion animals per year for food. We breed, raise, and kill more than 100 million animals per

year for research. We breed, raise, and kill countless animals each year for clothing, companionship, entertainment, and more. We capture or kill trillions of aquatic and land animals every year for food and other purposes. And of course, we are unsustainably consuming land, water, and energy, producing waste, pollution, and greenhouse gas emissions, and building a global infrastructure that excludes, harms, and kills other animals. As a result, we are destroying animal habitats, harming or killing trillions of animals, driving species to extinction, and causing ecosystems to collapse.

Our complicity in nonhuman suffering and death has clear moral implications. For a consequentialist, we have a moral obligation to improve animal welfare in the world because we have the power to improve animal welfare in the world. If you can save a deer drowning in the pond without sacrificing anything morally significant, you should. The same can be true for many other animals as well. For a non-consequentialist, we have a moral obligation to improve animal welfare in the world because we have a moral obligation to reduce or repair unnecessary harms that we cause. If you intentionally or foreseeably participated in pushing a deer into a pond, and if you can save the deer without sacrificing anything morally significant, you should. Once again, the same can be true for many other animals as well.

7. The Zoopolis. Generally speaking, humans can share agency more expansively than other animals can. Nonhuman animals can share agency and create implicit social contracts, with each other as well as with us. But only humans can share agency create explicit social contracts at the scale that we do. For better or worse, we have the power to jointly construct shared social, political, and economic systems that span nations and generations. As with all powers, we can, and do, use this power for both good and harm.

Many of the above harms result from our using this power for selfish purposes. But we can use them for altruistic purposes as well. If we create shared systems that promote human and nonhuman welfare simultaneously, then we can both promote good outcomes and address historical harms at the same time.

With great power comes great responsibility. For better or worse, we now have more power than any other species by far. We have a duty to exercise this power responsibly, by building shared systems that promote welfare and respect rights for humans and nonhumans alike. We might think that we can accomplish this aim by extending human political frameworks to nonhumans, for example by treating domesticated animals as citizens and wild animals as sovereigns.¹⁶ Or, we might think that we can accomplish this aim only by replacing human political frameworks with new multi-species political frameworks.¹⁷ There are other questions about political standing that we need to resolve as well. Either way, we have a collective responsibility to consider all animals, not only humans, when making decisions that affect them.

7. A moderate hybrid decision procedure

There are many reasons why a monist moral theory might imply a decision procedure that involves more Kantian duties in some contexts and more utilitarian duties in others. For example, many people agree that we might have reason limit the freedom of children more than the freedom of adults, that we might have reason to conduct harm-benefit analysis more when resources are scarce than when resources are abundant, and that we

¹⁶ For an example of such a view, see Donaldson and Kymlicka 2011

¹⁷ For an example of such a view, see Gabardi 2017.

might have reason to harm or kill people more in wartime than in peacetime. The reason is not that our fundamental values change in these contexts, but rather than the practical implications of our fundamental values change in these contexts. I will now suggest that the divide between humans and other animals can be similar, in part because it combines versions of these other divides together.

As a reminder, two-level utilitarians and rights theorists, of the kinds described above, can agree about a lot in practice. We have moral duties to humans and nonhumans all over the world. In order to treat humans and nonhumans well, we should generally respect rights, cultivate virtues, and cultivate relationships of care. Additionally, we should accept that morality is demanding, but we should also create space for our own projects and relationships. Similarly, we should accept that morality is restrictive, but we should also allow for exceptions; for instance, we can sometimes harm or kill individuals for their own good, in self-defense, in other-defense, or for the sufficiently greater good. We might accept this framework for different reasons, and we might develop the details differently, but we can agree on this much.

This decision procedure, combined with the differences considered above, might have a wide range of implications. First, there are several reasons why our duties to many humans might be more Kantian than our duties to many nonhumans. For a clear case, consider humans and ants. As we have seen, individual humans have much stronger interests than individual ants, as well as a much higher proportion of interests related to agency. Additionally, we are generally better able to interact with humans as individuals. Insofar as this is all true, we should think of human ethics more in terms of respect for individual rights by default. We should assign a higher value to human lives and a higher

bar for violating human rights. And of course, similar differences will apply, to greater or lesser degrees, across other species as well.

Second, and relatedly, there are several reasons why our duties to many nonhumans might be more utilitarian than our duties to many humans. For a clear case, consider once again humans and ants. Even if *individual* ants have weak interests, *populations* of ants can still have strong interests, since they can have so many members. Additionally, even if we are generally unable to interact with ants as *individuals*, we can still interact with them, and attempt to harm them less and help them more, as *populations*. Insofar as this is all true, we should think of ant ethics more in terms of the goal of aggregate welfare by default. Our focus should be less on improving individual lives and more on improving aggregate welfare. Once again, similar differences will apply, to greater or lesser degrees, across other species as well.

Third, insofar as a multi-species community involves more need and conflict, and less infrastructure for addressing this need and conflict, than a mostly human community, we will need to engage in more consequentialist reasoning. There will be more cases in which we need to use harm-benefit analysis to make tragic choices, as well as in which we need to restrain, harm, or kill individuals for their own good, in self-defense, in other-defense, or for the sufficiently greater good. Thus, whereas a mostly human society might allow for a more “Kantian” set of norms, with more individual liberty and fewer rights violations (not counting all the animals such a society excludes!), a multi-species society might require a more “utilitarian” set of norms, with more harm-benefit analysis and more rights violations.

Fourth, humans and nonhumans will play different roles in shared communities. On one hand, humans are moral and political agents. We have rights as well as responsibilities, including the rights and responsibilities of contractors. Additionally, we can motivate fellow humans to follow shared laws and treat themselves and others well through both rational and non-rational means. In contrast, nonhuman animals are moral and political patients. They have rights but no responsibilities, or, at least, they do not have the rights and responsibilities of contractors. Additionally, we are less able to motivate nonhumans to follow shared laws or treat themselves or others well through rational means, which means that we might have to rely more on non-rational means, such as priming, conditioning, or physical restraint.

Fifth, as we have seen, we might need to make a distinction between human and nonhuman rights for other, more contingent reasons as well. For example, insofar as there are more animals than we can support and we can support some more effectively than others, we might need to apply a discount rate across species, much as we do across nations and generations. Similarly, insofar as we need to take care of ourselves in order to be able to take care of others, we might need to create a moderate Kingdom of Ends for humans (and our companions) in order to promote animal welfare effectively more generally. Again, this might not be ideal. But this point of a decision procedure is not to tell us how to act in an ideal world, but rather to tell us how to act in the real world. And the real world is currently very far away from any reasonable ideal.

If we accept some combination of these ideas, then the result is a variety of norms across species. Near one end of the spectrum, we should accept a relatively Kantian set of norms for humans, by thinking of duties to humans primarily in terms of respecting

relatively strong individual agency rights. Near the other end of the spectrum, we should accept a relatively utilitarian set of norms for, say, r-selected wild animals, by thinking of duties to these animals primarily in terms of respecting relatively strong aggregate welfare rights. Other species will occupy other points along the spectrum, depending on a wide range of factors, including the strength and content of their interests and rights, our complicity in the harms that they endure or impose on others, and our ability to support them in treating themselves and others well effectively and sustainably.

However, it is important to emphasize that this moderate hybrid view is different from the traditional Nozickian hybrid view in many ways. First, this hybrid view is not based on a single binary difference, but is rather based on a series of overlapping differences, many of which come in degrees. As a result, this hybrid view will have different implications for different animals in different contexts. While our duties to humans will tend to be more Kantian than our duties to, say, mice, our duties to mice will tend to be more Kantian than our duties to, say, ants. And so on. For each individual, it all depends on what this individual is like, what our relationship with this individual is like, and what the broader social, political, and environmental context of our interactions with this individual is like.

Second, and relatedly, this view does not imply a purely Kantian set of norms for humans, nor does it imply a purely utilitarian set of norms for, say, r-selected wild animals (much less for other animals). While our duties to other humans might be relatively Kantian, they will also have consequentialist features, which might be more prominent in some contexts (such as wartime) and less in others (such as peacetime). Similarly, while our duties to, say, r-selected wild animals might be relatively utilitarian,

they will also have Kantian features, which might be more prominent in some contexts (such as when we live with a mouse) and less in others (such as when we consider how to treat mice in social policy). In all cases, the question is how regularly we can expect to face choice situations that require harm-benefit analysis and aggregation.

Third, and also relatedly, this view does not imply absolute priority of humans over nonhumans. For example, suppose that we can save either a human or 100 mice. In this case, we might have reason to save the human all else equal, since the human might have more at stake than all 100 mice. But now suppose that we can save either a human or 100,000 mice. In this case, we might have reason to save the mice all else equal, since they might have more at stake than the human. Moreover, even if we could permissibly save a human instead of 100,000 mice, that does not mean that we could permissibly kill 100,000 mice to save a human. The moral bar for harming others as means is higher than the moral bar for not helping others. Thus, our treatment of animals in food, research, and so on will still likely be morally wrong on this view.

Fourth, and relatedly, the implications of this view are contingent on facts about human psychology. Since this moderate hybrid view is meant to be a decision procedure, we need to strike a balance between (a) making it simple enough to guide our choices in practice, and (b) making it complex enough to capture the many differences that can obtain within and across species. For some purposes, such as creating public policy, we might have enough time and information to make many fine-grained distinctions. For other purposes, such as making everyday decisions, we might not, and so we might need to make fewer and more course-grained distinctions. Either way, plausibly our decision

procedure will need to be simpler than reality, but more complex than “Kantianism for people, utilitarianism for animals.”

Fifth, many of the implications of this hybrid view are also contingent on current social and political structures. For example, insofar as we learn more about animal welfare, reduce the activities that harm other animals, and increase the activities that support other animals, in part by building a more inclusive shared infrastructure, we might find that humans and nonhumans can co-exist peacefully much more than they currently can. This might increase the number of animals for whom relatively Kantian norms can be appropriate by default, and it might also increase the degree to which relatively Kantian norms can be appropriate for nonhumans more generally. Granted, we might not be able to change the world enough to eliminate these practical differences. But we can still change the world enough to reduce them.

So, even if a common moral standard applies to people and other animals in theory, we can and should still recognize a need to take a more Kantian approach to our interactions with some animals, such as humans, a more utilitarian approach to our interactions with other animals, such as insects, and various combinations of these approaches for other animals, depending on what these animals are like, how we relate to them, and a variety of other such factors. Again, this is not to say that we should take a purely Kantian approach to our interactions with humans or a purely utilitarian approach to our interactions with ants, much less with other nonhumans. But we do need different norms to ensure that they can live as freely as possible within the structures that allow us all to live and let live as much as possible.

It is important to emphasize that this moderate hybrid decision procedure is much more pro-nonhuman than the traditional Nozickian hybrid view. If we accept the kind of decision procedure that I have sketched here, then we should hold that our current policy of using trillions of animals each year for food, research, entertainment, and so on is morally wrong in practice. We should also hold that our current policy of destroying wild animal habitats and exterminating wild animals as a first resort whenever human and nonhuman interests seem to conflict is morally wrong in practice as well. So, to accept this decision procedure is not at all to vindicate our current anthropocentric social, political, or economic systems. It is instead simply to reaffirm the idea that equal consideration is compatible with differential treatment.

Granted, some people might find the relatively pro-human Nozickian hybrid view more intuitively plausible than this relatively pro-nonhuman hybrid view. But, first, the Nozickian hybrid view is less plausible at the theoretical level of morality, since it has less simplicity, power, and comparability. Second, the intuitive plausibility of this view is easy to explain away. We see this view as plausible partly because we have speciesist beliefs about welfare, agency, and morality. We see humans as more rational than we are and nonhumans as less rational (and sentient!) than they are. We also see humans as having the right to more than we do and nonhumans as having the right to less than they do, since we live in a society that normalizes human privilege and nonhuman oppression. Correct for these biases, and the intuitive case for the Nozickian view fades.

8. Conclusion

This discussion illustrates why a two-level view with a monist criterion of rightness and a hybrid decision procedure is an attractive alternative to a one-level monist or hybrid view. We start with single, unified view at the theoretical level (such as utilitarianism or rights theory, on a particular interpretation), and we then derive multiple contextual sets of norms at the practical level, including a moderate “Kantianism for people, utilitarianism for animals,” with many caveats and additional complexities. As a result, we can explain why we should accept multiple, seemingly incomparable sets of goods, duties, and virtues in practice, as well as why these sets of goods, duties, and virtues are unified and comparable in theory (as well as sometimes comparable in practice). Single-level monist and hybrid views, in contrast, lack this balance of virtues.

I think that this discussion supports a couple of more general, methodological conclusions as well. First, when we evaluate moral views, we should clearly distinguish the theoretical and practical levels of morality, and we should evaluate moral views according to different standards at each level. In particular, when we evaluate views at the theoretical level, we should expect them to have theoretical virtues such as simplicity, but we should not expect them to have practical virtues such as applicability. And when we evaluate views at the practical level, we should expect the opposite. Thus, for instance, it would be a mistake to dismiss monist views at the theoretical level for being too abstract and general, and it would also be a mistake to dismiss hybrid views at the practical level for not being unified enough.

Second, and relatedly, when we evaluate moral views, we might find that they converge more than we expected in practice. For example, I have suggested how utilitarianism and rights theory (on a particular interpretation) might at least partly

converge. On one hand, utilitarianism is less demanding and more restrictive in practice than we might have expected, owing to our psychological limitations, and it also involves more emphasis on rights. Meanwhile, rights theory is more demanding and less restrictive in practice than we might have expected, owing to a variety of factors, and it also involves more emphasis on harm-benefit analysis. While the details and reasons might be different, at least some other criteria of rightness (monist as well as hybrid) might converge on this kind of hybrid decision procedure as well.

To be clear, I have not argued for, or against, any particular moral view in this paper. Much more work is needed to assess how different theories work in practice, in a wide range of contexts, including but not limited to contexts involving human and nonhuman animals. Additionally, much more work is needed to compare sufficiently fleshed out monist views with similarly fleshed out hybrid views. With that said, I expect that, when the dust settles, a two-level theory with a monist criterion of rightness and a hybrid decision procedure will be more plausible overall than a one-level monist or hybrid view. I also expect that, generally and relatively speaking, this theory will imply more Kantian norms for humans and more utilitarian norms for many nonhumans, with plenty of caveats and additional complexities.¹⁸

References

Aristotle, 2016, *Nicomachean Ethics*. Scotts Valley: CreateSpace Independent Publishing

¹⁸ Thanks to David Killoren and Richard Rowland for editing this special issue, and for organizing the Oxford Workshop on Utilitarian Approaches to Animal Ethics in September 2019, at which I presented related material. Thanks also to the participants at this workshop for helpful feedback, and to Dale Jamieson for helpful discussion.

Platform

Bentham, Jeremy, 2018, *The Principles of Morals and Legislation*. New York: Franklin Classics.

Brink, David, 1989, *Moral Realism and the Foundations of Ethics*. New York: Cambridge University Press.

Chang, Ruth, 1997, *Incommensurability, Incomparability, and Practical Reason*. Cambridge: Harvard University Press.

Gabardi, Wayne, 2017, *The Next Social Contract*. Philadelphia: Temple University Press.

Hare, R. M., 1981 *Moral Thinking: Its Levels, Methods, and Point*. Oxford: Oxford University Press.

John, Tyler, and Jeff Sebo, forthcoming, "Consequentialism and Nonhuman Animals," in Douglas Portmore, ed. *The Oxford Handbook of Consequentialism*. Oxford: Oxford University Press.

Kagan, Shelly, 2019, *How to Count Animals, More or Less*. Oxford: Oxford University Press.

Kant, Immanuel, 2012, *Groundwork for the Metaphysics of Morals*. Cambridge: Cambridge University Press.

Korsgaard, Christine, 1996, *The Sources of Normativity*. Cambridge: Cambridge University Press.

Mackie, J. L., 1977, *Ethics: Inventing Right and Wrong*: 83-104. London: Penguin Books.

Nozick, Robert, 1974, *Anarchy, State, and Utopia*. New York: Basic Books.

Parfit, Derek, 2013, *On What Matters: Volume One*. Oxford: Oxford University Press.

----- 1984, *Reasons and Persons*. Oxford: Oxford University Press.

Railton, Peter, 1984, "Alienation, Consequentialism, and the Demands of Morality,"

Philosophy & Public Affairs 13:2: 134-171.

Ross. W. D., 1988, *The Right and The Good*. New York: Hackett Pub Co Inc.

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

Scheffler, 1982, *The Rejection of Consequentialism*, Oxford: Oxford University Press.

Sebo, Jeff, forthcoming, *Animals, Pandemics, and Climate Change*. Oxford: Oxford University Press.

Sebo, Jeff, 2015, "Utilitarianism, Multiplicity, and Liberalism," *Utilitas* 27:3: 326-346.

Sidgwick, Henry, 2014. *The Methods of Ethics*. New York: The MacMillan Company.

Swanton, Christine, 2005. *Virtue Ethics: A Pluralistic View*. Oxford: Oxford University Press.

Velleman, J. David. 2009, *How We Get Along*. Oxford: Oxford University Press.