

Commonsense Morality and Contact with Value

Adam Lovett and Stefan Riedener

Abstract. On the face of it, there are many kinds of moral duty. We should keep our promises; we should pay our debts of gratitude; we should compensate those we've wronged; we should avoid doing or intending harm; we should help those in need. These constitute, some worry, an unconnected heap of duties: the moral realm, or at least that of commonsense morality, is a disorganized mess. In this paper, we provide a unified account of commonsense moral duties. All these duties can, we think, be understood in terms of contact with value. You are in contact with a value when you are manifest in it or when it is manifest in you. When you're in contact with a value the weight of the reasons it grounds are amplified for you. And additionally, you have reason to get in contact with the good and avoid contact with evil. These ideas, we argue, can bring order to the chaos of commonsense morality.

Keywords. Gratitude · compensation · promises · doing and allowing · intending and foreseeing

1. Introduction

Intuitively, there are many domains of morality. Beneficence is one such domain: we should generally try to help other people. When a child is drowning in front of you, you should try to save them. But, on the face of it, beneficence is not the only moral domain. We should also keep our promises. When you promise a friend to pick them up from the airport, you should do so, even if it would do more good to spend the time working for effective charities. Equally, we should pay our debts of gratitude. When someone does a good thing for you, you should do good things for them, even if you would make the world better by helping the global poor. Conversely, we should compensate people when we wrong them. If you break someone's leg, you should at least pay for their medical bills, even though the money would surely do more good spent on malaria nets. And we shouldn't do, or intend to do, harm. Indeed, it is worse to do harm than to merely allow it and it is worse to intend harm than to merely foresee it. Intuitively, morality is pluralistic: there are many different kinds of moral reasons and they do not simply reduce to one another.¹

This is, of course, contrary to utilitarianism. Utilitarianism is the monistic moral theory *par excellence*. It says that, despite appearances, morality does all reduce to beneficence: it is entirely about making people's lives better. But most philosophers reject

¹ The outlined view is Ross's (2002 [1930]) view.

utilitarianism.² At root, its problem seems to be that it is an impersonal or impartial or agent-neutral moral theory. As Henry Sidgwick said, for utilitarianism, the moral point of view is the “point of view of the universe” (1981 [1907], 382). But you are not the universe. You are your own person, with your own point of view. You should keep your promises even when it would be impersonally better to break them; you should not be impartial between helping your benefactors or victims and helping complete strangers; you should not be neutral between punching someone and letting someone be punched. As W.D. Ross put it, “the essential defect of [utilitarianism] is that it ignores [...] the highly personal character of duty” (2002 [1930], 22). The personal character of duty is the thread that runs through the moral domains besides beneficence. It has not been satisfactorily woven into the utilitarian tapestry. This, we suspect, is why many philosophers are receptive to the pluralistic view of morality.

Yet, there is something quite unsatisfying with this pluralistic view. It seems to make morality an “unconnected heap of duties.”³ It robs us of any general explanation of what is right and wrong, and instead provides us with an inelegant, disunified collection of claims about the moral sphere. This points to the great strength of utilitarianism: it’s a stunningly simple, unified and elegant moral theory. Simplicity, unity and elegance are all general criteria for theory-building: they contribute to the appeal of theories in any domain. This, we think, is responsible for the magnetic attraction of utilitarianism and the notoriety of pluralistic deontology. The aim in this paper is to rectify this problem: it is to present a unified theory of non-utilitarian ethics. More precisely, we aim to provide a unified account of compensation, gratitude, promises and non-maleficence. W.D. Ross thought that these exhausted the non-utilitarian parts of morality. Perhaps that is false; perhaps there are other domains of non-utilitarian ethics. Yet these are extremely important domains, so a unified account of these domains is of great import in itself.⁴

Our view is that our moral reasons can be understood in terms of contact with value. Intuitively, you have very different relationships to different good and bad things. For instance, you’re in close contact with the suffering of a person you’ve harmed. Their suffering manifests your recklessness or ill will. You don’t have the same kind of relationship to the suffering of people whom you’ve never affected. Similarly, you’re in close contact with the virtue of those who benefit you. Their good will is manifest in your experiences. You don’t have the same kind of relationship to the virtue of strangers on the other side of the planet. Our core idea is that these differing relationships, this contact with value, shapes the landscape of your reasons. It does so in two ways. The first way is *reactive*: you have especially weighty reason to react to values with which you’re already in contact. The second way is *proactive*: you have reason to get into contact with the good and avoid being in contact with the bad. We call this the Contact Account of commonsense morality (sections 2-3). We think it can

² According to Bourget and Chalmers (2021), fewer than a third of philosophers are consequentialists.

³ For the term, and a critical discussion of the charge, see McNaughton (1996).

⁴ There are, of course, other attempts to provide unified, non-utilitarian accounts of such domains. Kantians try to reduce them to the Categorical Imperative (Kant 2011 [1785]). Contractualists reduce them to principles we could not reasonably reject (Scanlon 1998). Yet each view is some distance from commanding universal consent, so we think it worthwhile to present a different theory.

provide a unified understanding of the personal character of duty. Or at least, we'll argue, it can provide a unified account of our reasons of compensation (section 4), gratitude (section 5), promise-keeping (section 6) and non-maleficence (section 7).

Let's be explicit about our methodology. We think a good account of gratitude, say, will explain how gratitude works. Various principles govern our reasons of gratitude. For example, you have weightier reasons to those who benefit you supererogatorily than those who have a duty to help you. A good account of gratitude will explain the truth of these principles. The same is true of compensation, promises and non-maleficence. Now of course, it's contentious what principles govern any moral domain. Some people, for example, deny that you owe more gratitude for supererogatory beneficence. But we take it that, *prima facie*, this is plausible. A view that explains such appearances is better than one that does not. More fundamentally, the appropriate methodology on these issues is akin to that of reflective equilibrium. Principles about particular moral domains should be judged in part by how well they fit into a wider ethical framework, and a wider ethical framework should be judged in part by its ability to generate plausible principles. This can require adjustment on both sides: we can adjust specific claims to fit into a broader framework, and a broader framework to get us specific claims. Ultimately, it is the coherence and intuitive plausibility of the resultant picture that redounds to its credit.⁵ This grounds the appeal of the Contact Account of commonsense morality. Let's now spell out its core notions.

2. Contact with value

We'll start by clarifying what 'value' means to us. We agree with utilitarians that the world is full of valuable things. But we think it's not just well-being and ill-being that are valuable. Consider a beautiful painting or a loving relationship. Both ground certain reasons. You have reason to appreciate the painting and to protect the relationship. Consider, also, a piercing pain or an unjust institution. These also ground such reasons. You have reason to alleviate the pain or to rectify the injustice. The reasons, in these cases, are both practical and affective: they are reasons to do certain things and have certain emotions. You should both try to alleviate someone's pain and have an emotional response to their being in pain: you should regret it, you should feel sorry for them. And these reasons are, in a certain sense, noninstrumental. You don't have reason to alleviate pain due to the distant causal consequences of such alleviation; you have reason to do so because pain is bad in itself. Most generally, we'll say that a *value fact* is a fact that grounds a noninstrumental practical or affective reason. As these examples suggest, value facts can be both good and bad. Good value facts, such as the fact that a relationship is loving, warrant positive responses. Bad value facts, such as the fact that an institution is unjust, warrant negative responses. Generally, anything which grounds some such reason is a value in our sense.

Next, let's clarify the relation of 'contact'. To convey a sense of this notion, consider the idea that if you already are in contact with value, this intensifies your reasons with respect to it. This idea seems intuitive, even outside of morality. Take art. You are in

⁵ This methodological view is most strongly associated with Rawls (1999).

closer contact with a beautiful painting you look at every day than with one you've never really laid eyes on. This gives you weightier reason to appreciate the former painting, or to protect it if it's in danger. Here your contact with the painting's beauty is passive: it impinges itself on you. Similarly, suppose you write an eloquent book. You are in closer contact with your book than you are with someone else's. This gives you weightier reason to try to publish your book than similar books by others, or to regret the manuscript's rejection by a publisher. Here your contact with your book's eloquence is active: it flows from your own capacities. Both sorts of contact seem to intensify the weight of your reasons grounded in the relevant values.


Let us pursue a more precise understanding of 'contact'. A straightforward idea is that contact is a causal notion: you are in contact with the things that causally affect you and with the things that you causally affect. You're in contact with the painting's beauty because your aesthetic experience is a causal product of that beauty. You're in contact with your book's eloquence because it is the causal product of your own actions. This idea is simple and elegant; tenable, unfortunately, it is not. You are causally connected with many things to which you lack a normatively relevant sort of contact. Imagine that a thousand years ago two of your ancestors were transfixed by the beauty of some church. They met while admiring it, fell in love, and a thousand years later you were born. The beauty of the church was a key causal contributor to your existence: had it not been for it, you would never have come to be. But that does not intensify your reasons grounded in this beauty. You don't need to feel any special aesthetic response to that church. A simple causal account of contact, then, seems implausible. Causal connections are often too contingent or accidental to be fit for our purposes.

We need a less accidental connection between you and value. We propose to understand this connection in terms of manifestation. Manifestation is primarily understood in terms of dispositions. Consider salt's solubility. The solubility of some salt is its disposition to dissolve when you put it in water. Here its being in water is the *stimulus condition* of the disposition and its dissolving is its *manifestation condition*. The former is the thing that activates, or stimulates, the latter.⁶ Thus, when you put salt in water, and it dissolves, the dissolution manifests the solubility. When you drink the water and throw up, that does not manifest the solubility of salt. It is merely a causal consequence of it. Or consider the fragility of a vase. When you drop a vase and it shatters, the shattering manifests the vase's fragility. When the shopkeeper scolds you for breaking their vase, this is caused by the fragility but does not manifest it. People's dispositions can also be manifest. Michelangelo was disposed to produce beautiful artworks when given the time and resources: the beauty of the *Pietà* manifests this disposition. Lyndon B. Johnson (LBJ) was disposed to bully those he had power over: him berating his subordinates manifests that. Gandhi was disposed to eschew violence even in the face of deep provocation: the nonviolence of the Indian independence movement manifests this disposition. Paradigmatically, then, manifestation is the non-accidental connection between dispositions and their manifestation conditions. We will say that you're in contact with a value when your dispositions are manifest in that value, or when that value's dispositions are manifest in you.

⁶ For an overview on the metaphysics of dispositions, see Choi and Fara (2018).

Let us say a bit more about manifestation. Here we want to distinguish a few different notions. We start with actualization. If D is a disposition with stimulus condition S and manifestation condition M , we'll say that M *actualizes* D when M occurs because D and S obtain. In this sense, the salt's dissolving actualizes its solubility when it dissolves because it's soluble and it is in water. Plausibly, manifestation is a little more than mere actualization. To see this, consider again the beauty of a painting. One might think this beauty is not itself the disposition to cause aesthetic experiences in people that look at the painting. The beauty and this disposition are not identical. Yet, still, those aesthetic experiences do, in a perfectly natural sense, manifest the painting's beauty. To capture this, we now introduce the notion of ground. Grounding is a notion of non-causal explanation. It is the connection between crimson and red, between the members of a set and the set as a whole, or between the parts of a table and the table.⁷ Plausibly, the painting's beauty grounds its disposition to cause aesthetic experiences when looked at: the existence of the disposition is non-causally explained by the beauty. Roughly, we want to say that manifestation is indifferent to such connections of ground. More precisely, let's say that p is *ground-theoretically connected* to q if and only if p grounds q or q grounds p . We say that q *manifests* p if and only if q or something ground-theoretically connected to q actualizes p or something ground-theoretically connected to p . This definition picks out connections just as non-accidental as actualization. It precisely articulates our conception of manifestation.

It will help to illustrate this notion diagrammatically. In the following diagram, upward arrows stand for the grounding relation and the horizontal arrow stands for actualization:



Here t grounds disposition D , D grounds s ; v grounds q , and q grounds u ; and q actualizes D . So u , q and v manifest s , D and t . Thus, for example, imagine you look at a beautiful painting and have an aesthetic experience. The painting is disposed to cause aesthetic experiences (D) and this disposition is grounded by (rather than identical with) the painting's beauty (t). Thus, when that disposition is actualized by your aesthetic experiences (q), these experiences manifest the beauty. Alternatively, imagine you cruelly berate someone, causing them pain. Here we can think of your cruelty (D) as a disposition that is actualized in your victim's pain (q). Suppose, additionally, that the pain fits into a more general story of your victim's oppression (u): they are often bullied. Then your cruelty is manifested in their more general oppression. This identifies the notion of manifestation that puts you in contact with a value. Heuristically, it is almost always unproblematic to think of manifestation as the simple connection between a disposition and its actualization. But our official, more involved,

⁷ The relevant notion of ground is strict partial ground. For more on it, see e.g. Fine (2012).

characterization of manifestation invokes ground-theoretic connections. Contact with value flows through this notion of manifestation.

We've spelt out what 'value' and what 'contact' is, but what is it for 'you' to come into contact with a value? Who are you, anyway? In answering this question, we don't aim to provide an account of who you are that is fit for all purposes. We don't, for example, aim to contribute to the literature on personal identity or the metaphysics of persons with our account of 'you'. We merely aim to provide a quasi-technical account that helps underwrite the extensional plausibility of the Contact Account. In this sense, we take you to be the sum total of your experiences, attitudes, actions and dispositions. You are what you think, feel, see, desire, do or are disposed to do. One can think of this as a collection of facts. There are facts about what attitudes, actions and dispositions you have or have had: these facts together make up who you are. That you raised a child, love cats or once felt the pain of a broken toe are each part of what makes you you. That you were adventurous in your youth and will be pusillanimous in old age are part of who you are. For our purposes, you are the collection of facts about your experiences, attitudes, actions and dispositions. When one of these facts makes contact with a value, we say that *you* are in contact with that value.

Your contact with a value, in sum, involves three elements: the two relata—'you' and the 'value'—and the relevant relation itself—the relation of 'contact'. Let's clarify one final thing about this relation. In our view, contact comes in degrees: you can be in more or less contact with a value. This is because both actualization and grounding come in degrees. On degreed actualization, think about when a vase shatters because you dropped it versus when it shatters because you dropped a tank on it; the former is more an actualization of its fragility than the latter. This itself should likely be understood in terms of causal contribution: the more of a causal contribution a disposition makes towards its own manifestation condition, the more that manifestation condition actualizes the disposition. On degreed grounding, consider the magnificence of the *The Godfather* films. This is grounded to a greater degree in Al Pacino's performance than in Sofia Coppola's; the former is more responsible for the magnificence than the latter. This allows us to define a notion of centrality. A fact is more central to you insofar as it makes a bigger ground-theoretic contribution to who you are; a fact is more central to a value insofar as it makes a bigger ground-theoretic contribution to the value. Our view is that you are more closely connected to a value when things more central to you are manifest to a greater degree in things more central to the value, or vice versa. Thus the extent to which you are in contact with a value is a function of the magnitude of certain actualization and grounding relationships.

3. The contact principles

We can now state the two principles that make up the Contact Account of commonsense morality. Let p be some value fact. Our first principle is the

Reactive Contact Principle: The weight of the reasons grounded by p are amplified for you to the extent that you are in contact with p .

This principle says that your contact with a value increases the weight of the noninstrumental practical or affective reasons this value grounds: you have weightier

reason to respond to values you're in contact with than to respond to equal values with which you lack such contact. So, for example, you have more reason to protect the painting you've appreciated many times than a painting you barely know. And you have more reason to regret the rejection of your manuscript than the rejection of someone else's. In both cases, your contact with a value amplifies the weight of the reasons it grounds.

Our second principle is the

Proactive Contact Principle: If p is good, then you have reason to put yourself in contact with p , and if p is bad, then you have reason to avoid being in contact with p .

This principle says that you have reason to put yourself in touch with the good and isolate yourself from evil. So, for example, suppose you haven't yet looked at this beautiful painting or written your eloquent book. Both the painting and your book are (or would be) good. Looking at the painting would put you in contact with the former good, and writing your book would put you in contact with the latter. Thus, by the proactive principle, you have a reason to look at the painting and write your book. In both cases, your actions can put yourself in contact with the good.

How are these principles connected? We think they're metaphysically independent, in the sense that neither grounds the other. But there are logical connections between them. To show this, we need two further principles. On the one hand, we need a principle linking reasons for attitudes to those for action. This says that if, were you to do something, you would have reason of a specific weight to be glad (regret) that you did it, then you actually have reason of such weight (not) to do that thing.⁸ On the other hand, we need a transmission principle: that, if you have reason of some weight to be glad about (regret) something, then you have reason of that weight to be glad about (regret) anything that made it the case you have such reason. With these principles in hand, the Reactive Contact Principle implies the Proactive Contact Principle. Here's how. When something is good, you have reason to be glad about it. When something is bad, you have reason to regret it. By the Reactive Contact Principle, when you get in contact with a value, these reasons are amplified. So, by the transmission principle, were you to put yourself in contact with a good (bad), you would have weighty reason to be glad about (regret) having put yourself in contact with that good (bad). But then, by the linkage between actions and attitudes, you actually have weighty reason to put yourself in contact with goods and avoid putting yourself in contact with bads. And that is the Proactive Contact Principle. So the contact principles form a logically coherent picture of normative reality (although neither grounds the other).

This connection allows us to illuminate the relative weight of the different reasons the contact principles generate. We think that the weight of your proactive reasons (not) to come into contact with p correspond to the weight of the reactive reasons you *would* have to respond to p *were* you in contact with p in that way. Let's illustrate this idea with our examples. Suppose you've written your book. By the Reactive Principle, you have

⁸ There's some similarity between this principle and so-called 'reflection principles' (Arntzenius 2008).

a reason to desire to publish it: such a desire is a fitting response to its value. In addition, suppose you contemplate destroying that beautiful painting. By the Proactive Principle, you have a reason not to do that: doing so would put you in contact with the bad of the painting's destruction. Now, if you *were* to destroy that painting, you would have a reactive reason to desire that you hadn't done that. And suppose that, in that case, it would be fitting for your desire not to have destroyed the painting to be equally as strong as your desire to publish your manuscript. Then your proactive reason not to destroy that painting is equally as strong as your reactive reason to publish your manuscript. So proactive and reactive reasons are comparable in weight.

These principles will together make up a cohesive account of commonsense ethics. Now one might wonder *why* each principle is true. We're inclined to take them to be fundamental. Every conception of morality (Utilitarianism included) must reach explanatory bedrock somewhere: we suspect the contact principles are that bedrock. Of course, one might still challenge them. Perhaps they rely on too pluralistic a picture of values, or they suggest that harms done out of character are not wrong, or they imply you should avoid learning about evil. We'll address these objections in section 8. But, for now, let us turn to our argument for these principles: they provide an elegant explanation of how compensation, gratitude, promises and non-maleficence work.

4. Compensation

We'll start with a vignette. Imagine you are driving home after a long day. It is dark and wet. You make a right turn. You hear a bump, and your heart hits the roof of your mouth. You hadn't properly checked your mirror, and you've knocked over a cyclist. You stop quickly and get out of the car. They're injured. You call an ambulance. It later becomes clear that you've broken their leg. In this case, you should compensate the cyclist: you should make their life better because of the harm you've inflicted on them. Specifically, you have stronger reason to help this particular cyclist than to help an injured cyclist you had nothing to do with. Additionally, your compensation should fit their harm. It would be odd to just buy them a vacation in Hawaii or finance the couples therapy they're undergoing with their partner. It's their leg you've broken, so that's what you should fix. You should pay for their medical bills. A slight modification of the vignette illustrates a third principle about compensation. Suppose that hitting the cyclist was not accidental: you did it in order to knock them out of an upcoming cycling race, and help a good friend of yours win it instead. Your friend does win the competition, and with it \$10,000. Here, it's not just you, but also your friend who has special reason for compensating the cyclist. When someone benefits from the wrongdoing of others, they should compensate the person wronged.⁹

Our aim is to explain these three features of compensation. Let's start with the first and most basic point: when you harm someone, you have special reason to compensate them. You have weightier reason to help them than to help someone whom you haven't harmed. This derives from the fact that, when you harm someone, that

⁹ For more about compensation, see e.g. Walker (2006) or Radzik (2009). For the idea of compensatory duties after benefiting from someone's wrongdoing, see e.g. Butt (2007) or Goodin and Barry (2014).

generally puts you in contact with a bad. In the above vignette, your carelessness consisted, in part, in the disposition to knock over cyclists when turning. That disposition is manifest in the cyclist's broken leg. So this puts you in contact with the bad of that injury. The same point applies very broadly when you harm someone. Suppose your anger causes you to break someone's nose, your financial irresponsibility causes someone to go bankrupt, or your callousness causes emotional hurt. In each of these cases your dispositions are manifest in some sort of injury: a physical, financial or emotional injury. And that puts you in contact with these bads. Yet when something is bad, we have reason to get rid of it. The badness of a broken leg gives us reason to help the person who suffers it. And so, according to the Reactive Contact Principle, you have an especially weighty reason to help the people you've harmed. This reason is weightier than your reason to help people that were harmed in other ways, because you're in closer contact with the injuries you inflicted yourself.

Now, this doesn't mean you always have very strong reasons to compensate those whom you've harmed. Consider, for instance, a harm to which your victim has consented. Suppose you fought a professional boxing match, and broke your opponent's nose. Then, you still seem to have some special reasons to help them. If you have a choice, say, between helping your opponent and helping another professional boxer who got equally injured in another match by someone else, it seems you should favor your victim. However, your reasons of compensation are much weaker than if your victim never consented to being punched. The Contact Account can explain this as follows. Generally, harms to which the victim consented are less impersonally bad than harms to which they didn't consent. Suppose you're an uninvolved third party, and come across two people with a broken nose: the first is a professional boxer, who had consented to fighting a match; the other simply became the victim of an unconsented assault. Intuitively, you have stronger reasons to help the latter. According to the Reactive Contact Principle, when you are manifest in a value fact, the reasons it grounds are amplified for you. Thus, you'll have stronger reasons to compensate your victim if you broke their nose without their consent than if you broke it in a fight to which they agreed.

Let's turn to the second aspect of compensation: that compensation should fit the harm. Again, if the cyclist is also having troubles with their spouse, it wouldn't really be appropriate to help them fix their marriage rather than their leg. Reasons of compensation are not generalized reasons to improve the lives of people you've harmed. They're reasons to get rid of the specific bads you've put into their lives. The explanation of this is simply that it's these specific bads that you're in contact with. You're not in contact with the cyclist's marital troubles. These troubles don't manifest any of your dispositions. It's the broken leg that does. And so, by the Reactive Contact Principle, you have special reason to pay for the cyclists' medical bills, but not to help them with their marriage. The point generalizes: when you harm someone, you put yourself in contact with a specific bad. It's that bad you have special reason to erase, and in this sense compensation should fit the harm.

We now look at the third aspect of compensation. You have special reasons to compensate people who were wronged if you benefitted from that wrongdoing. Again, the Reactive Contact Principle can explain this. Consider your friend, who won the

competition because you broke that cyclist's leg. Here, part of what grounds the badness of the broken leg is the cyclist's inability to compete in that competition. This inability is a disposition: the disposition, among other things, to remain at home when the race starts, and to have to let other people win it. This disposition is manifest in the victory of your friend. In other words, the badness of the injury of that cyclist is manifest in the victory of your friend. So your friend has a special reason to care about this badness: to alleviate it or pay for the cyclist's medical care. This point also generalizes. When you benefit from a wrongdoing, very often your benefit manifests a harm done to whoever was wronged. This puts you in contact with that harm, and thus gives you special reason to ameliorate it. So, when you benefit from a wrongdoing, very often you should compensate the person wronged.

The Reactive Contact Principle, thus, helps us explain three core features of compensation. Our discussion here does, however, raise a question. You don't merely have weighty reason to compensate those you harm: you have a duty (or obligation) to do so. You owe them compensation. Yet our contact principles only give you weighty reasons to compensate. How can we get duties out of that? This is a problem for anyone who starts ethical theorizing with talk of reasons; one needs some way of deriving duties from reasons.¹⁰ But this problem has a straightforward solution: you have a duty to do something when you have weightier moral reason to do it than you have reason not to do it. You have a duty to compensate somebody when your moral reason to compensate them outweighs your reason not to compensate them. Generally, duties should be understood in terms of the relative weight of your reasons. That also applies to gratitude, promise-keeping and non-maleficence. So let us turn to how the Contact Account can illuminate gratitude.

5. Gratitude

Imagine you're walking home late at night, and someone starts hassling you. You're worried about your safety: you think your unwelcome interlocutor could turn violent at any minute. A stranger sees the situation and walks over. They're concerned about your well-being and so they ask you if you're being bothered. The intervention dissuades your interlocutor: they walk away. You should be grateful to the stranger. You should feel especially positively towards them, and you should be especially disposed to do them good in return, if you can. If you see them in trouble in some future night, you should walk over and try to help them. This illustrates the core feature of gratitude: when someone benefits you benevolently, out of concern for your well-being, you have reason to be grateful to them. If the stranger had benefited you out of self-interest, merely because they wanted you to owe them a favor, you would have no debt of gratitude to them.¹¹ Notice, further, that being benefitted supererogatorily amplifies the weight of your reasons of gratitude. If the stranger went beyond the call of duty in benefiting you, if they had no obligation at all to intervene in your situation,

¹⁰ For a recent exploration and defense of "reason-first" views, see esp. Schroeder (2021).

¹¹ For this kind of claim, see Berger (1975).

you should show them more gratitude. In this section, we will explain these features of gratitude.

Let's start with the core feature: that benevolently bestowed benefits generate debts of gratitude. Benevolence consists in a concern for people's well-being. Such a concern consists, in part, in the disposition to help people out when you see that they are in trouble. Thus, when the stranger helps you, your experience of being helped manifests these dispositions. So that puts you in contact with their benevolence. The key further idea is that benevolence grounds certain reasons. Specifically, benevolent people deserve good things. We have weighty reason to help out benevolent people and to feel positively towards their life going well. To see this, think about how you should respond to two lives: a good life lived by a benevolent person, and a good life lived by someone who never acts to help others. You should be more pleased by the first life, and more disposed to bring it about. By the Reactive Contact Principle, when you're in contact with someone's benevolence, the weight of these reasons is amplified. Thus, when someone benefits you out of benevolence, you have especially weighty reason to want their life to go well and to help them out. Thus we have reasons of gratitude when we're benefited benevolently. We can contrast this with a self-interested motivation. Suppose someone benefits you just because they want you to benefit them in return. This puts you in contact with this person's self-interest. But that someone is self-interested doesn't ground reasons to care about them. So we can explain the core feature of gratitude: being benefited specifically out of benevolence gives us reason to be grateful to our benefactor.

We now turn to the fact that you should show more gratitude to benefactors who went beyond their duty in helping you. We explain this similarly. When someone benefits you supererogatorily, they have a disposition to do more than morality requires of them. They show, at least when benevolently motivated, a small measure of moral saintliness. This disposition is manifest in your experience of being benefited. Thus, you are in contact with their moral saintliness. But moral saints deserve good things. Compare two lives: a good life lived by someone who often acts supererogatorily, and a good life lived by someone who never does more than they morally need to. You should be more pleased by the first life than the second, and more disposed to bring it about. Thus, we have reason to help those who are disposed to go beyond the call of duty, at least when they are appropriately motivated. So by the Reactive Contact Principle, when someone benevolently benefits you supererogatorily, you have especially weighty reason to make their life go well. The weight of your general reason to help out the morally saintly is increased by your contact with your benefactor's specific saintliness. The Contact Account, then, can explain how benevolence and supererogation matter to gratitude.

The Contact Account also illuminates a third aspect of gratitude. We've been discussing what's sometimes called prepositional gratitude. This is gratitude *to* someone. But there is also a notion of *propositional* gratitude. You can simply be grateful that something is the case rather than grateful *to* someone. Suppose you've spent many days immersed in the beauty of the Grand Canyon. In this case, we shouldn't say you're grateful *to* the Canyon for being so beautiful. Rather, you're grateful *that* it is. Prepositional and propositional gratitude are distinct. Yet, they seem to have

something in common: we at least use the same word to refer to both, and phenomenologically they seem similar. But what do they have in common?¹² The answer, we think, is that both involve you being in contact with something good. To be immersed in the Grand Canyon's beauty is to be in contact with that beauty. By the Reactive Contact Principle, such contact gives you especially weighty reasons to appreciate and protect that beauty. This is similar to your especially weighty reason to protect your benefactors: both derive from the fact that non-accidentally benefiting from something puts you in contact with its value. So, the Contact Account can explain the connection between prepositional and propositional gratitude.

It appears to us, then, that the Reactive Contact Principle sheds substantial light on how gratitude works. It is perhaps worth making a further point at this juncture. In all our explanations we're relying on claims about values, dispositions, grounding or manifestation: about what values and dispositions there are, what grounds them, and about what manifests what. We think all of these claims are independently plausible. But we also think they draw plausibility from the coherence of the overall picture. The fact that certain tenable claims about manifestation, for example, help us explain the nature of gratitude gives us some evidence for these claims. These claims are mutable points in the dynamics of reflective equilibrium.

6. Promises

Suppose you promise to pick up a friend from the airport. Having made the promise gives you special reason to do so. What explains this? One might think that this brooks no explanation: it's a basic, fundamental fact that we should keep our promises. But that seems peculiar: intuitively, we should be able to explain promissory reasons in some way. And if we're able to explain gratitude and compensation in more fundamental terms, this intuition seems doubly strong: it seems especially peculiar to think promissory reasons are more basic than these other domains. Moreover, it seems plausible that we should be able to explain promissory reasons in terms of interests. Keeping your promises serves people's interests, and this is why you should do so. Tellingly, most of the literature on promising articulates some version of this view. Some people think that the interest underpinning promissory reasons is an interest in having our expectations satisfied.¹³ Others think that it's an interest in managing our relationships.¹⁴ And some think that it's an interest in not acting out of line with our joint intentions.¹⁵ So, this basic intuition is widespread: promissory obligations brook an explanation in terms of the interests of the promisee.

Unfortunately, this intuition is difficult to vindicate. The key issue is that me having an important interest in your doing something simply doesn't imply that you have a weighty reason to do it. This is for two reasons. First, not all of my interests give you weighty reasons. Suppose I have a weighty interest in your dropping out of a job

¹² For discussion of this distinction, see McAleer (2012) and Manela (2015).

¹³ For different versions of this view, see Thomson (1990), Scanlon (1998, ch.7) and Kolodny and Wallace (2003).

¹⁴ For this sort of view, see Shiffrin (2008).

¹⁵ This is a slight twist on the view in Gilbert (2011).

interview, since you're my competition. This doesn't give you a weighty reason to drop out. Second, although I might have an important interest in your doing something, other people might have more important interests in you not doing it. These contrary interests have to be weighed against my own, and beneficence will generally direct you to do what promotes people's interests most overall. So it's unclear how to vindicate the intuition that promises can be explained in terms of interests. Why does the promisee's interest give you a weighty reason, or a weightier reason than similar interests by others?

The Reactive Contact Principle can explain this. At an abstract level, the crucial thought is that, when you make a promise to someone, you put yourself into contact with the special interests underpinning promises. This interest, whatever it is, grounds reasons: specifically, it grounds reasons to do the things you said you would do. Thus, you have special reason to keep your promises. Let's illustrate this with some concrete theories. Consider the view that the interest underpinning promissory obligations is that in having expectations satisfied. When you promise someone that you'll do something, you typically raise their expectations that you'll do it. Their raised expectations are a manifestation of your abilities: specifically, the ability to make people promises. They wouldn't have the expectations if you hadn't made the promise. But, on the view under discussion, that someone expects you to do something gives you a reason to do it. Usually, this would be a relatively lightweight reason. But, when the expectation manifests your abilities, by the Reactive Contact Principle that amplifies the weight of this reason. Thus, you have weighty reason to keep your promises. Similarly, consider relational theories. These say that when you break a promise to someone, you damage your relationship with them, and their interest in preserving that relationship underpins promissory reasons. The fact that it would damage your relationship not to do what you promised is a manifestation of your ability to make commitments: it wouldn't damage the relationship had you not made a promise. Additionally, this fact about relational damage grounds a reason to do as you promised. Thus, given the Reactive Contact Principle, when you've promised someone something you have especially weighty reason to do it. So both these theories illustrate the explanatory force of the Contact Account. Let's be clear; we're not endorsing either theory. We just use them to showcase how the Contact Account can explain why promissory obligations are underpinned by interests.

There are two further features of promising that the Contact Account sheds light on. First, promises give rise to secondary obligations. Suppose you break your promise to pick your friend up from the airport. This doesn't free you from any special reasons towards your friend. At a minimum, you should tell them you won't be able to pick them up and apologize for breaking your promise. More expansively, you have special reason to avoid them being harmed by your promise-breaking: to pay for their taxi, say.¹⁶ This is all because you remain in contact with the promise-underpinning interest even when you've broken your promise. You remain in contact with people's interest in being able to manage their relationships, have their expectations satisfied or whatever that interest is. But this interest still generates reasons after the promise-

¹⁶ See e.g. Riedener and Schwind (2022).

breaking. Specifically, you have reason to repair the damage represented by the frustration of this interest. So you have especially weighty reason to repair the harm that you caused. Consider how this works on the expectational view. When you jilt your friend, how bad the frustration of their expectations is depends on the cost to them of getting home from the airport. By mitigating this cost, through paying for their taxi, you mitigate the harm that you've caused. This gives you secondary reasons even when you've acted against your primary reasons to keep the promise.

Second, some promises are weightier than others. Promises to your nearest and dearest are weightier than promises to complete strangers. Other things equal, it's worse to break a promise to your mother than to a salesperson. Also, promises whose violation causes more harm are weightier than promises whose violation is more harmless. It's worse to jilt your friend at the airport when this means that they have to walk for two hours than when it just means they must take a twenty-minute bus. The former fact holds because we're in closer contact with the interests of our friends and family members than with those of strangers. That includes whatever general principles underpin the fact that we have interests implicated in promising. For example, perhaps it is important to have our expectations satisfied because we're autonomous. Well, you're in closer contact with the autonomy of your nearest and dearest than with the autonomy of perfect strangers. So your promises to friends and family members are weightier. The second of these facts holds because the weightier is the interest you're in contact with, the weightier will be the resulting reason. The weight of this interest is in part modulated by how serious its frustration is, and so that seriousness impacts the weight of a promise.

In sum, then, the Reactive Contact Principle illuminates why interests underpin promises, how promises give rise to secondary reasons, and which promises are weightier than others. Now, we don't claim to have explained every feature of promises here, or indeed every feature of the other domains we discuss. For example, promises seem invalid when coerced, or if the promisee doesn't accept them. We haven't tried to explain that in this section. But we think the success the Contact Account has in explaining the features of promising we've focused on is reason enough to think it plays an important role in understanding promises.

7. Non-Maleficence

Imagine you and three strangers have been kidnapped. The kidnapper gives you a choice: you can either shoot one stranger, or they will shoot the other two. This is a difficult decision. It seems like there is a moral difference between doing harm and merely allowing it to be done. Specifically, we usually have weightier reason to avoid doing a harm than to avoid allowing it.¹⁷ Equally, imagine that you're a bomber pilot in a just war. It may well be permissible to bomb a munitions factory with the aim of disrupting wartime production, even if you foresee that ten civilians will die as a result of the bombing. But it would not be permissible to bomb a housing estate with the

¹⁷ For this point, see Williams (1973). For more recent discussions, see e.g. Scheffler (2004) or Woollard (2015).

intention of killing ten civilians in order to damage the enemy's morale, even if the consequences for the enemy's fighting effectiveness were the same. There is a moral difference between intending and foreseeing. We usually have weightier reason to avoid bringing about intended harm than to avoid bringing about merely foreseen harm.¹⁸ These are well-known distinctions. In this section, we show that the normative significance of both can be explained in terms of contact with value.

We'll start with two observations. First, every harm takes a quite specific form. When someone is shot, their death is replete with concrete details: they died in a particular time and a particular manner. The bullet entered their body at a certain angle, it interrupted their bodily functionings in a certain way. When you shoot someone, more of your dispositions are manifest in these details than when you merely allow someone to be shot. You determine the exact time and manner of their death. These detailed facts ground the fact that they died. So, when you shoot someone, you're more in contact with their death than when you merely allow them to be shot. More generally, you're typically more in contact with harms you do than those you merely allow. Second, intentions are just bundles of practical dispositions. To intend to kill ten people is, in part, to be disposed to bring about the death of ten people. Foreseeing, in contrast, is a purely epistemic notion: to foresee something is to accurately believe that it will happen, but it needn't involve any disposition to bring about what is foreseen. Thus, when you intentionally bring about a harm, *ceteris paribus*, it manifests more of your dispositions than when you merely foresee a harm resulting from your actions. So, standardly, you are in more intimate contact with harms you intend than those you merely foresee.

We now use these observations to explain the distinctions. Here we rely on the Proactive Contact Principle: the principle that says you have reason to get in contact with good things and avoid contact with bad ones. As we've just seen, you are in closer contact with harms you do than those you merely allow, and with those you intend than those you merely foresee. By the Proactive Contact Principle, you have reason to minimize your contact with bads. Doing harm puts you in closer contact with a bad than merely allowing it, and thus you have weightier reason to avoid doing harm than to avoid allowing it. Likewise, intending harms puts you in closer contact with a bad than merely foreseeing it, and so you have weightier reason to avoid intending harm than merely foreseeing it. And that is just what we sought to explain. The Contact Account gives us straightforward explanations of why harms we do are typically worse than those we merely allow and why those we intend are typically worse than those we merely foresee. It illuminates the distinctions in maleficence.

Let us be clear just what we have explained. We've explained the moral difference between certain paradigm cases of doing versus allowing and intending versus foreseeing. To focus on the former, we've explained why it's usually worse to shoot someone than to let someone be shot. We haven't given an explanation of why it would *always* be worse to do harm than allow harm. We're not sure it is always worse to do harm than to allow harm. Suppose, for example, you involuntarily spasm and cause some harm, although you could have stopped yourself from spasming. Perhaps this is

¹⁸ For classic discussions of this distinction, see Foot (1967) or Quinn (1989).

an allowing. Yet it is not clear that our evaluation of this harming depends on whether it is a doing or an allowing. The distinction, in this case, doesn't seem very important. So we think the thing to be explained is not the normative significance of these distinctions in full generality, but rather their significance in certain paradigm cases. That is what we have explained.

Nonetheless, our account is consistent with reasonable intuitions about more *recherché* cases. Consider, for example, the removal of barriers. Suppose someone you despise is hiking at the bottom of a large ravine. If you push a boulder down the ravine, this will do harm to them. But suppose you remove a log from the path of a boulder that is already rolling down the ravine. Plausibly, this is just an allowing of harm. One might think the former is worse than the latter.¹⁹ To explain this, we need some claims about causal contributions. Specifically, we need to claim that your dispositions make more of a causal contribution to a harm when you initiate a harmful sequence (by pushing the boulder) versus when you simply remove a barrier to that sequence. The idea justifying this claim would be that pushing a boulder is a full cause of the picknicker's death, but the absence of a barrier to the boulder is a mere enabling condition for death. Enabling conditions, one might think, make a lesser causal contribution to what they enable than do full-blown causes. So you are in closer contact with the death when you push the boulder than when you remove the barrier. Thus, the Contact Account implies that the former is worse than the latter. We are not certain that these claims about causal contributions are true.²⁰ But if one denies them, it also seems implausible that removing a barrier is any worse than initiating a harm. So the Contact Account explains the pattern of reasonable judgments about such cases.

Let's make a final, related, point. We've emphasized that, sometimes, the difference between doings and allowings might not be normatively significant. The same is certainly true of that between intendings and foreseeings. To see this, imagine that you're stuck in a cave and a man is wedged in the exit. He is blocking your escape. You blow him into pieces with dynamite, clearing your way out. You say you only intended to blow the man into pieces, and merely foresaw that this would lead to his death. This isn't a sound defense of your actions. The man's death and you blowing him up are too closely connected for the distinction between intending and foreseeing to matter here. But when, exactly, are two outcomes too closely connected for the distinction to have normative traction?²¹ The Contact Account answers this question. On our account, when p grounds q , p and q are too close for the intending and foreseeing distinction to have moral traction. In such a case, even if you merely foresaw q , bringing it about is just as bad as if you intended it. To see why this is, notice that the man's body being blown apart partially grounds his death. So, when you intended to

¹⁹ For these judgments, see Woollard (2015, 71–75). We're neutral on them.

²⁰ A reason to deny these claims is that they gel poorly with prominent metaphysical theories of causation. If you think causation is counterfactual dependence, for example, it's unclear why removing a barrier makes less of a contribution than initiating a sequence. For a discussion of such theories, see e.g. Menzies and Beebe (2019). One of us thinks "so much the worse for these theories"; the other is inclined to be more deferent to metaphysical theory. But, overall, this conflict makes us less confident in the causal claims.

²¹ This question vexed both Foot (1967) and Quinn (1989).

blow the man up, you're in contact with his death. So by the Proactive Contact Principle, your relationship to his death is, normatively speaking, the same as your relationship to something you straightforwardly intended. You antecedently had weighty reason not to blow the man up. This is a virtue of the Contact Account: it illuminates when the intending/foreseeing distinction is normatively (in)significant.

8. Objections and Replies

We've laid out an abductive argument for the Contact Account. We now want to explore and address some objections to the account. The first objection targets this argument. We've said that we've provided a simple, unified accounts of a large swathe of non-consequentialist ethics. Yet we invoke a plurality of values: we've claimed that beauty, pain, injury, benevolence, moral saintliness and much else besides are values. One might worry, then, that we have not really unified anything at all. Our axiology, one might fret, is deeply disunified. In reply to this worry, we agree that our account of values is disunified. But we don't think that is a very serious problem. That's because everyone's picture of values should be disunified. Remember that a value is just something that grounds a reason. Intuitively there are many different things that ground reasons. The view that our reasons to feel awe in the face of the Grand Canyon and those to help children in need are grounded in the exact same kind of value is simply implausible. Reasons for awe or admiration or promotion or protection are often grounded in different kinds of values. So every extensionally adequate ethical theory should accept a disunified axiology.²² Thus, it is no special cost of our theory that it too recognizes a plurality of values. And nor does it mean the Contact Account has not unified anything. It has not unified the realm of value, but it has unified the realm of reasons. The pluralistic conception of morality holds that there are multiple entirely distinct domains of reasons. The Contact Account shows that all these domains can be reduced, in one way or another, to contact with value.

A further, related, worry concerns our framing. We've said we're not utilitarians. Yet we talk a lot about values for non-utilitarians. We explain the normative significance of compensation and gratitude and promises and so on ultimately in terms of what is good and bad. So one might worry, are we not closet utilitarians after all? We are not. The key difference between the Contact Account and utilitarianism is the agent-relativity of the Contact Account. Utilitarianism denies that the agent's relationship to any value shapes their reasons: our view is built around the idea that it does. This is how our view captures the highly personal character of duty. Now there are highly heterodox versions of consequentialism that are also agent-relative. These views take what is valuable to be relative to each agent: what is good from your perspective is different from what is good from my perspective.²³ Yet the Contact Account doesn't collapse into these views either. These views say that the good matters ultimately because we have reasons to promote it. Our view, in contrast, recognizes many reasons

²² Indeed, paradigmatic deontologists like Ross (2001[1930]), and most contemporary consequentialists (e.g. Railton 1984, 148–5; Brink 1989, 211–90) even accept a pluralistic picture of practical values: the kind of values that ground reasons for actions. Axiological pluralism is the norm, not the exception.

²³ See Portmore (2011, 84–117).

beyond those to promote something: it recognizes reasons to admire, to honor, to protect the good too. On both the Contact Account and any version of consequentialism, values play a critical role in generating practical reasons. But consequentialism is monomaniacally focused on promotion; the Contact Account admits of many other ways to appropriately relate to values.

Let's look at a third, more concrete, concern about the Contact Account. Sometimes, one does bad things that are out of character. Imagine that you are a kind, gentle person. But you have a bad day, and you hurt someone needlessly. Their having been hurt doesn't manifest your deep character traits. Can we say, still their hurt has a special normative significance for you? That you antecedently had special reason not to hurt them and now have special reason to compensate them? We can. The key point is that your simple-but-fleeting dispositions coincide with more complicated, more stable dispositions. You might not generally have a disposition to hurt people needlessly. You wouldn't hurt people needlessly had you had a good day. But the fact that you did hurt someone needlessly means you have a disposition to do so when a very particular stimulus is in place: when the particular facets of your bad day are realized. This disposition is not so fleeting, and it manifested when you hurt someone. So you are in contact with their pain. Perhaps the simpler dispositions play a bigger role in grounding who you are because they are manifest more often, and so when you hurt someone in an in-character manner perhaps those hurts have more normative significance for you. But you are still in contact with the values manifest by seemingly fleeting dispositions.

Let's turn to a final criticism of the Contact Account. This critique swings on certain kinds of passive contact with the bad. By the Proactive Contact Principle, you have reason to avoid such contact. Yet you can, plausibly, acquire such contact merely by finding out what is going on in the world. Reading about bad events is, plausibly, a way to get into contact with those events. Yet, intuitively, we shouldn't avoid such passive contact with the bad. We shouldn't shut ourselves off from all the evil in the world. So, one might think, this aspect of the Contact Account is doubtful.

There is a straightforward reply to this worry. Although you have reason to shut yourselves off from the world, you have countervailing, usually weightier, reason to inform yourself about the world. This countervailing reason comes from two sources. On the one hand, it comes from the fact that willful ignorance, and especially willful moral ignorance, is a vice. It is a bad character trait to stick your head in the sand. Epistemically isolating yourself puts you in contact with this vice. So you have reason not to epistemically isolate yourself. On the other hand, the countervailing reason comes from the fact that you have reason to put yourself in contact with the good. When you help people that puts you in contact with weighty goods. When you save someone's life you are in contact with the good of their continued existence. When you give someone deworming pills, you're in contact with the good of their renewed energy. So you have reason to effectively help people. You cannot do that if you isolate yourself from the world's evils. Attempts to help, when done out of ignorance, are typically inefficacious. So you have reason to inform yourself. Thus, on balance, you should not shut yourself off from all that is bad in the world; you have weightier reason to learn about the world's evils than you do to isolate yourself from them entirely.

That address the most obvious challenges to the Contact Account. There are of course other objections to the account we might address. But one can reply to the most salient concerns about our picture.

9. Conclusion

We have spelt out a unified account of compensation, gratitude, promises and non-maleficence. On our view, all these moral domains arise from the normative import of contact with value. And this account, we've shown, can illuminate various principles about how each domain works. That, to reiterate, provides abductive evidence for the account. The contact principles draw their main support from their explanatory power: they bring together a vast swathe of disparate moral phenomena into one ethical framework and draw plausibility from doing so. In this way, the Contact Account has some of the virtues of utilitarianism; it provides a unified, elegant, account of what we ought to do. Yet it avoids Utilitarianism's vices: it can explain the highly personal character of moral duty. The Contact Account provides us with a unified, general and satisfactory account of commonsense morality.

It may be obvious to the reader that that is not all it can do. Our contact principles, and especially the Proactive Contact Principle, can also illuminate beneficence. We mentioned this in the prior section: we have reason to help other people because doing so puts us in contact with the good. Thus, the Contact Account can assimilate beneficence to the rest of morality; it gives us an explanation of the home ground of utilitarianism. In other work, we argue that contact with value also illuminates the normative import of special relationships, discounting and personal well-being. It explains why we should care more about our friends and family than about strangers, or more about temporally nearby events than events in the far future or past, and it furnishes us with an account of prudential reasons.²⁴ So, we think that contact with value might provide a quite comprehensive account of the normative domain. Yet spelling this out in depth is a project for another time.

References

- Arntzenius, Frank. "No Regrets, Or: Edith Piaf Revamps Decision Theory." *Erkenntnis* 68, no. 2 (2008): 277–97. <https://doi.org/10.1007/s10670-007-9084-8>.
- Berger, Fred R. "Gratitude." *Ethics* 85, no. 4 (1975): 298–309. <https://doi.org/10.1086/291969>.
- Brink, David Owen. 1989. *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Butt, Daniel. "On Benefiting from Injustice." *Canadian Journal of Philosophy* 37, no. 1 (2007): 129–152. <https://doi.org/10.1353/cjp.2007.0010>

²⁴ See [Redacted] and [Redacted].

- Choi, Sungho, and Michael Fara. 2018. "Dispositions." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.
- Fine, Kit. "Guide To Ground." In *Metaphysical Grounding: Understanding The Structure of Reality*, edited by Fabrice Correia and Benjamin Schnieder. Cambridge: Cambridge University Press, 2012.
- Foot, Philippa. "The Problem of Abortion and the Doctrine of Double Effect." *Oxford Review* 5 (1967): 5–15.
- Gilbert, Margaret. "Three Dogmas about Promising." In *Promises and Agreements*. Oxford University Press, 2011.
- Goodin, Robert E. and Christian Barry. "Benefiting from the Wrongdoing of Others." *Journal of Applied Philosophy* 31, no. 2 (2014): 363–376. <https://doi.org/10.1111/japp.12077>
- Kant, Immanuel. *Groundwork of the Metaphysics of Morals*, edited and translated by Mary Gregor and Jens Timmermann. Cambridge: Cambridge University Press, 2011 [1785].
- Kolodny, Niko, and R. Jay Wallace. "Promises and Practices Revisited." *Philosophy & Public Affairs* 31, no. 2 (2003): 119–54.
- Manela, Tony. "Gratitude and Appreciation." *American Philosophical Quarterly* 53, no. 3 (2016): 281–94.
- Menzies, Peter, and Helen Beebe. 2020. "Counterfactual Theories of Causation." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2020. Metaphysics Research Lab, Stanford University.
- McAleer, Sean. "Propositional Gratitude." *American Philosophical Quarterly* 49, no. 1 (2012): 55–66.
- McNaughton, David. "An Unconnected Heap of Duties?" *The Philosophical Quarterly* 46, no. 185 (1996): 433–47. <https://doi.org/10.2307/2956354>.
- Quinn, Warren S. "Actions, Intentions, and Consequences: The Doctrine of Double Effect." *Philosophy and Public Affairs* 18, no. 4 (1989): 334–351.
- Radzik, Linda. *Making Amends: Atonement in Morality, Law, and Politics*. New York: Oxford University Press, 2009.
- Railton, Peter. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13, no. 2 (1984): 134–71.
- Riedener, Stefan and Philipp Schwind. "The Point of Promises." *Ethics* 132, no. 3 (2022): 621–643. <https://doi.org/10.1086/718080>.
- Rawls, John. *A Theory of Justice. Revised Edition*. Cambridge MA: Belknap Press of Harvard University Press, 1999.
- Ross, W. D. *The Right And The Good*, edited by Philip Stratton-Lake. Oxford: Oxford University Press, 2002 [1930].
- Scanlon, Thomas. *What We Owe to Each Other*. Cambridge MA: Belknap Press of Harvard University Press, 1998.
- Scheffler, Samuel. "Doing and Allowing." *Ethics* 114, no. 2 (2004): 215–39. <https://doi.org/10.1086/379355>.
- Schroeder, Mark. *Reasons First*. Oxford: Oxford University Press, 2021.

- Shiffrin, Seana Valentine. "Promising, Intimate Relationships, and Conventionalism."
The Philosophical Review 117, no. 4 (2008): 481–524.
<https://doi.org/10.1215/00318108-2008-014>.
- Sidgwick, Henry. *The Methods of Ethics*. 7th ed. Indianapolis: Hackett, 1981 [1907].
- Smart, J. J. C., and Bernard Williams. *Utilitarianism: For and Against*. Cambridge:
Cambridge University Press, 1973.
- Thomson, Judith Jarvis. *The Realm of Rights*. Harvard University Press, 1990.
- Woollard, Fiona. 2015. *Doing and Allowing Harm*. Oxford: Oxford University Press.