

Victoria McGeer
and Philip Pettit
'The Empowering
Theory of Trust' in
Paul Faulkner and
Thomas Simpson,
eds, *The Philosophy
of Trust*, OUP

2

The Empowering Theory of Trust

Victoria McGeer and Philip Pettit

We have argued independently in earlier papers that if I trust you in a certain domain, then that may help to improve your capacity to prove reliable and deserving of trust; it may have a positive, empowering impact on your psychology (Pettit 1995; McGeer 2008b). It is a matter of common assumption that any capacity may come in degrees or levels; more on this later. But how can my relying on you in a trusting manner raise the degree to which you are reliable and deserving of trust? We sketch a defence of the possibility in this paper and elaborate on some of its implications.

The paper itself is in three sections. The first provides a background account of what it means for me to rely on you and trust you, and what it means for you to have the capacity to respond: to be trust-responsive. The second explores the case for the theory that my trusting you can increase your trust-responsiveness. The third looks at some implications of the fact that your trust-responsiveness is sensitive in that way to my trusting you.

1. Trust and Reliance

The most basic idea is that of reliance. If I am to count as relying on you, in our usage, then I must assume that you have a disposition to act in a certain manner and I must act on that assumption. Thus I may assume that you are disposed to put out your garbage for collection on the appropriate day and treat your putting it out as a signal that today is the day. Or I may assume that you have a disposition to check the weather report before you go out and treat your taking an umbrella as a prompt to prepare for rain. Or I may assume that you have a disposition to recognize what is best in the long-term for your children and be ready to follow the same line in raising my own. In any such case I display an attitude of reliance insofar as I assume the presence of a disposition in you that I treat as a potential guide for my own behaviour. And I perform an act of reliance in letting that assumption dictate how I do indeed behave on this or that occasion.¹

¹ In all of these cases you act in a way appropriate to the purpose on hand: to getting the garbage collected, or to guarding against rain, or to raising your children. But I might rely on you in a certain sense so

If this is what relying on you involves, what extra is needed for it to be the case that I trust you? In ordinary usage, the word ‘trust’ does not always come apart from the word ‘rely’. But in the terminology we use in this paper, and in keeping with a well-established philosophical tradition, it does come apart (Baier 1986; Holton 1994; Jones 1996). ‘Trust’ in that tradition of thinking is a term of art.

In this usage, I will count as trusting you to do something X just insofar as three conditions are met. First, I manifestly rely on you to do X: I make clear to you my assumption that you will prove reliable in doing X.² Second, I assume that the manifest fact of my reliance will weigh with you as a reason for choosing voluntarily to X. And third, this assumption helps to explain or reinforce my relying on you.

What does it mean to say that the manifest fact of my reliance weighs with you as a reason for choosing voluntarily to X: that is, to do X without being forced to do it? That fact must provide you with a reason, by your view of things, for proving reliable, whether by making X-ing as such attractive or by making it attractive as a means to some other end. In particular, it must provide you with a reason to X without making X-ing into a more or less involuntary choice: that is, without making not X-ing into an intuitively unacceptable alternative, as it would do if it meant, for example, that you would be punished for not X-ing (Olsaretti 2004). The reason for X-ing which the manifest fact of my reliance on you provides may actually motivate you; the motivation may be sufficient to get you to X or a necessary part of a condition sufficient to get you to X. Or, to take a further possibility, it may reinforce some independently sufficient motive you may have for X-ing: it will be there, if needed, to provide back-up assistance.

As reliance involves both an attitude and an action so too, on this account, does trust. I will display an attitude of trust just insofar as I assume that the manifest fact that I am relying on you will weigh with you as a reason to prove reliable; henceforth we take the requirement of voluntariness to be understood. And I will exercise an act of trust just insofar as I act on that assumption or belief.

On this usage, I may often display or exercise reliance without displaying or exercising trust. I regularly rely without trust on inanimate objects and non-human animals, as when I act as if the bridge I am crossing will not collapse, as if the weather tomorrow will be fine, or as if the dog will not soil the carpet. And I may even rely on people in a way that does not involve trust. When I go on a political demonstration, I rely on others to join me but I do not display or exercise trust. I do not expect that they will have been aware of my intention to participate—it may not have been manifest to them—let

far as you act in some regular way, even perhaps inappropriately by my lights to the situation at hand. Believing you always put out the garbage a day too early, I might rely on your doing so and put out my garbage the day after. In the first sort of case, the reliance is straight, as we may say, in the second bent. When we speak of reliance in the text, it is always straight reliance we have in mind.

² This condition may be weakened to agree with a developmental view of trust, according to which I may sensibly manifest reliance without having much confidence that you will actually prove reliable. In that sort of case, mentioned at the end of this section, I manifest the belief in your reliability with the aim of enhancing your capacity to prove reliable over time.

alone aware of my relying on them to take part as well. And when I rely on some rough-necks not to assault me in a public place, this reliance will not typically involve trust either. While I may take them to recognize that I am relying on them in this way, I am unlikely to think that that will weigh with them as a reason to prove reliable. I will probably think that the only consideration that moves them is the fear that an assault would prompt third parties to come to my aid.

Suppose, by contrast with such cases, that I have been sitting near you in a restaurant and ask you to keep an eye on my computer when I go to the toilet. I am manifestly relying on you not to make off with the computer, and not to let anyone else do so either. And to the extent that I rely on you in this way, taking you to be significantly responsive to that reliance—taking you to have the capacity to be moved by my relying on you—I count as putting my trust in you. I have and reveal an attitude of trust in putting myself in your hands—this may be based on how you look or on how you seemed in an earlier exchange of pleasantries—and in acting on that attitude I exercise trust: I perform an act of trust.

In offering this account of trust, we take your responsiveness to my manifest reliance—your trust-responsiveness, for short—as a capacity: specifically, as the capacity to be moved voluntarily to action by the manifest fact of my reliance. You will have the capacity involved just insofar as you satisfy an important modal condition. For an open range of variations on the actual situation that keep in place the fact that I am manifestly relying on you, you would be moved to action in most of them; or at least you would be moved to action when those variations did not introduce countervailing reasons, impeding passions, or anything of a recognizably disruptive kind.³ In short, the capacity consists in a disposition to be robustly, if not inevitably, moved to action by the fact that I am manifestly relying on you.⁴

By all accounts a capacity like trust-responsiveness may come in degrees and this will be important in later discussion. Taking the capacity as a disposition of the kind described, it turns out that there are two ways in which trust-responsiveness, may vary in degree; it may vary in dependability, as we put it, or in durability. The capacity will be present with a higher rather than a lower degree of dependability to the extent that the disposition is relatively sensitive to appropriate triggers: it is fit to be activated over a greater rather than a lesser number of scenarios in which the relevant reasons for trust-responsiveness are in place. By contrast, the capacity will be present with a higher rather than a lower degree of durability to the extent that it is relatively resistant to extraneous features that inhibit its manifestation: it is fit to display itself over a greater rather than a lesser number of scenarios where potential disrupters are present (Pettit 2015, ch. 2).

³ The idea that you would be moved to action in most of the scenarios is not essential; it might be replaced, for example, by the idea that you would have a significant chance of being moved to action in all of them. See Pettit (2015, Appendix II).

⁴ For a characterization of capacity on these lines see Smith (2003) and McGeer and Pettit (2015).

To illustrate the distinction between dependability and durability in our usage, consider the mechanical disposition of a sophisticated thermostat designed to maintain ambient room temperature at a set degree. The system will be more or less dependably disposed to do this job just to the extent that it can cope with various, limited increases in cold or heat: they trigger it, as appropriate, to turn on the central heating or the air conditioner. But, however dependable, the system will not be very durably disposed to do its job if the power supply is erratic. And its durability will be increased, independently of any impact on its dependability, if the supply becomes less erratic or a back-up source of energy is put in place.

The same distinction between durability and dependability will apply with any human disposition or capacity that we ascribe, including that of trust-responsiveness, insofar as the following conditions obtain. First, the disposition is interpreted as a disposition to respond dependably to a certain range of triggering scenarios: this, in the way the heating/cooling system is interpreted as a mechanism for restoring the pre-set temperature in response to a limited range of increases in cold or heat. Second, there are some scenarios where the disposition fails to generate the expected response, by analogy with the scenarios where the heating system does not work. Third, the intuitively best explanation of this failure is that in those scenarios there are factors present that disrupt or mask the operation of the disposition, as the failure of the electric supply disrupts the heating system.

There are many cases where such conditions are satisfied, allowing us to distinguish quite naturally between the dependability and durability of various ordinary human capacities. We may be happy to ascribe an outstanding mathematical capacity to you, treating it as highly dependable—you can solve a great range of tricky problems—while recognizing that the ability is disrupted or masked by inhibition in the company of hostile or assertive competitors. We may be happy to ascribe a highly dependable capacity to play very difficult pieces on the piano—a capacity demonstrated time and time again—while admitting that, unfortunately, it is disrupted or masked by stage fright in the presence of a large audience. And we may be confident in treating you as a dependably kind person—this shows up in how you routinely treat your family and friends, colleagues, and strangers—while acknowledging that this capacity is disrupted or masked when you imbibe a few too many alcoholic drinks. By analogy with these cases, we hold that the trust-responsiveness that I ascribe in investing my trust in you may be suitably dependable—it may mean that you score quite highly among your peers—without always being very durable.⁵

⁵ While this is the possibility of particular interest in this paper, it is worth noting more generally that there is a sense in which dependability and durability may vary independently of each other, so that the disposition or capacity may not only be dependable without being durable, for example, but also durable without being dependable. If we map dependability on the horizontal axis of a graph, and durability on the vertical, it is possible for a disposition to be situated at any set of coordinates, disappearing from view only at 0,0.

Consider, for instance, the scenario that Plato makes vivid in the *Republic*, invoking the ring of Gyges. Though it may be true that your trust-responsiveness is highly dependable, it is reasonable to ask whether you would retain that disposition—or any other of a range of virtues—if you were provided with a ring that made your actions undetectable and that put you effectively beyond rebuke or penalty. Assuming you think that you probably would not retain the disposition in that scenario—that you would probably be corrupted by the ring—should you conclude that you are not as dependably trust-responsive or kind or honest now as you thought you were? Not if, as seems reasonable, you think that the dependability of those dispositions is measured only over scenarios where no magical intervention occurs and you live in the presence, and under the potential regard, of others. Not, in other words, if you interpret the disposition as a disposition to respond dependably to just those scenarios. What the ring of Gyges shows, on this interpretation, is that virtuous dispositions that we prize in ourselves and others, however dependable they are, may not be durable in the measure to which we thought them.⁶

The analysis of trust outlined in this section of our paper does not yet provide support for the empowering view of trust. It is quite consistent with the analysis, that when the manifestness of my reliance motivates you to prove reliable, that is entirely because of a pre-existing disposition on your part not to let people down when they are manifestly relying on you; that disposition may not be affected in any way by my trusting you. The act of manifesting reliance may serve as a trigger that activates that disposition, in other words, without giving it a boost. It may prompt you to reveal your established capacity to be motivated to act as others manifestly depend on you to act but not serve in any way to improve that capacity. The empowering theory goes further in suggesting that by manifestly relying on you—by exercising trust—I may not only cause you to exercise your existing capacity for trust-responsiveness; I may also cause you to develop that capacity, achieving a higher degree of dependability or durability.

Although the empowering theory of trust is more encompassing than the view implied in the bare analysis of the phenomenon, we take it here in a way that makes it less encompassing than a third approach. Our focus is on straightforward cases where I expect that the manifestation of my relying on you will generally ensure that you actually prove reliable; it will help to prompt or reinforce your acting as I rely on you to act. But the considerations that support the empowering theory, so understood, suggest that there will also be cases where it makes developmental sense for me to manifest reliance on you, even though I do not expect you actually to prove reliable: certainly I do not expect this with any degree of confidence. Manifesting reliance will make developmental sense to the extent that I believe that you can be brought to develop trust-responsiveness over time, as a result of my efforts; you may be an adolescent, to

⁶ Pettit (2015, ch. 2) argues that the situationist claims in recent psychology do not establish that virtues are undependable and in that sense rare—this is the standard interpretation—but only that they are not as durable as might have been thought; they are subject to unexpected disrupters, as various experiments in social psychology have shown. See also Merritt (2000).

take a salient case, and I may be a parent who is anxious to make you trust-responsive. The empowering theory might be taken as a theory governing a spectrum of cases that includes these examples as well as the straightforward cases, and that also includes examples that do not fall neatly into either category. But for simplicity we restrict attention here to the theory as it applies in the straightforward cases.

We try to sketch some considerations that support the theory in section two. And then in the third section we elaborate some implications of the theory for how we should think about trust-responsiveness and, more generally, about capacities of the kind it exemplifies. The considerations we will be introducing in support of the empowering theory are broadly empirical. But they are rooted in matters of common experience and assumption and we think of the case they make for the theory as an elaboration of common-sense tenets. Still, the case we present is certainly subject to testing and might be usefully brought into contact with the empirical literature on trust. Unfortunately we are not in a position to do that ourselves, and certainly not within the confines of this paper.

2. The Case for the Empowering Theory

2.1 *Two theories*

When I display or exercise trust in dealing with you, as we saw, then I act on the belief that you are responsive to my manifestly relying on you to act in a certain way. In taking this line I make myself vulnerable to you. Depending on our relationship with one another, and indeed on what is at stake, this may involve exposing myself to any of a variety of ways in which you may let me down (Baier 1986; Holton 1994; Jones 1996). It is one thing for me to trust you, a complete stranger, with my computer; it is quite another to trust you, a friend, with a secret. But whatever the mode of vulnerability involved, I trust you just insofar as I take my manifest reliance and my associated vulnerability as a reason that weighs with you to prove trust-responsive, whether that reason serves in an active or back-up role. If I did not believe that you were responsive to such considerations of reliance—for short, to reasons of trust—then I would not count as putting my trust in you.

But to be responsive to reasons of trust—to have the capacity to be suitably motivated by the manifest reliance of others—is an ambiguous notion. It may mean either or both of two different things.

First, it may mean that you have a standing disposition to be responsive to reasons of trust, where the reasons present may vary depending on your relationship to the trustee and even to what is at stake; you may have different ~~and stronger~~ reasons for being responsive to the trust of a friend from the reasons you have to be responsive to the trust of a stranger. But within a given category—say, in your relationship to a given friend—your sensitivity to the reasons themselves is a standing feature of your psychology.

Alternatively, to be responsive to reasons of trust may mean that you are ready or predisposed, on being presented with evidence of trust, to form a disposition there and then, in the situation at hand, to respond to reasons of trust. Where in the first example your responsiveness is a standing, independent disposition, in the second it is a disposition that is actually formed or developed in response to another's display of trust.

An analogy may help make this difference clear. Consider what it means to have an *adaptive* immune system. Creatures with such a system are disposed to fight off infections but they may be disposed to do this in either of two different senses. They may have been exposed to a non-lethal form of the disease, in which case they will have developed the relevant antibodies and will be actively disposed—disposed in a once-for-all-inputs fashion—to fight it off: the antibodies will be there, waiting to get to work when suitably triggered. But even if they have not been exposed to the disease and not developed this active immunity, the creatures will be disposed in a more remote, input-by-input sense to fight it off; by hypothesis, their adaptive immune system will start manufacturing antibodies in response to exposure. This passive immunity is not quite as effective as the active sort but it can provide a perfectly adequate level of protection.

We may all agree that when I trust you I must believe that you are responsive to reasons of trust: that is, to considerations to the effect that someone—say, someone in a particular relationship—is manifestly relying on you in a certain way. But you may be trust-responsive in either or both of two ways, just as your body may fight off disease in either of two ways. Setting aside the complication raised by the relativization of trust to different relationships, you may have a standing disposition to be motivated by reasons of trust across all instances of reliance. Or you may have an input-sensitive predisposition or readiness to develop the disposition to respond appropriately to reasons of trust in specific instances of reliance. Where the disposition you act on in the first case has a standing character, the disposition you act on in the second has a case-specific character: it is formed anew in each situation.⁷

Though not all trust-theorists subscribe to this view, the standard or received theory of trust, as we shall characterize it, holds that trust-responsiveness consists in a once-for-all-cases type of responsiveness. On this picture, trust-responsive people have a standing disposition, more or less constant across different situations, to respond to reasons of trust. If standing trust-responsiveness is equated with trustworthiness, in the ordinary sense of that term, then the claim of the received theory is that people's trust-responsiveness is nothing more or less than their (standing) trustworthiness. It is

⁷ Consistently with this distinction, it is possible that as I and others display trust in you and help to elicit a situational disposition to prove reliable, this will have a boosting effect on your standing disposition over time. The idea is that the more you are situationally activated to prove responsive to trust, the more dependable your standing disposition to prove responsive will become. It may even be that situational activation remains essential to keeping the standing disposition in place at a minimal level of dependability. This would be analogous to the immunity case, where it may be that while past exposure to disease has elicited antibodies that provide you with active rather than just passive immunity, continuing exposure—regular booster shots of a vaccine, for example—may be necessary to maintain that active immunity.

a virtue that they bring to every situation and provides a solid, unchanging base on which we can predict what they will do, or be inclined to do, now in this scenario, now in that.⁸

The empowering theory of trust does not deny that people may be trustworthy in this standing sense and that this disposition may play a role in helping to explain why they do not let down their trustors. But it adds that how trust-responsive such trustees are in dealing with another may also be a function of their situational sensitivity to the presence of that person in the role of a salient trustor: someone who is manifestly relying on them to behave in a certain manner. The idea is not that the trustor may have a special hold on them, although that may be true too. It is rather that any trustor is liable to elicit a novel degree of trust-responsiveness in a trustee—not necessarily the same degree for every trustor—as well as triggering the trustee's existing trustworthiness.

On the empowering theory, then, the net trust-responsiveness that you display, now in this scenario, now in that, may be a function of two variables. The first variable is how trustworthy you are in general: how much sensitivity you have to reasons of trust—reasons of manifest reliance—in themselves. The second is how far you are moved by the presence of those who manifestly rely on you in certain ways: how sensitive you are to the fact that here and now others are putting their trust in you. We will look shortly at how to model the contribution of each of these factors in representing your degree of trust-responsiveness. But first we must look at the precise claims of the empowering theory and the case for accepting them.

Suppose for a given scenario of trust—a case where someone manifestly relies on you—that you are trust-responsive with a certain degree of dependability and a certain degree of durability. You are disposed to respond to those reasons across a range of possibilities that measures the degree of dependability; and your disposition is proof against a range of disrupters that measures the degree of durability. It may be in such a case that your responsiveness reflects just a standing disposition, as in the received theory. Or it may be that in each relevant case it also reflects the effect of a case-specific disposition. Thus it may be that your final disposition is the resultant of standing and situational components.

The empowering theory is committed to two theses. First, that there are grounds for thinking that even if you bring a standing capacity of trust-responsiveness to a given situation of trust, you are likely to be influenced within that situation in a way that enhances the capacity; you are likely to develop a disposition with a higher degree of dependability in reaction to situational inputs. And second, that there are grounds, not

⁸ This take on trustworthiness is motivated only by considerations of convenience and clarity. The theory outlined in the paper might also be construed to show that trustworthiness is not a standing trait of the kind presented here; it might identify trustworthiness with trust-responsiveness. As we use the term here, 'trustworthiness' measures your standing responsiveness to reasons of trust, where as we have noted, the requirements of those reasons may vary in keeping with your relationship to the trustor (stranger, colleague, friend) and in keeping with what is at stake in the context (their computer, their secret, or whatever).

just to think that you can be made trust-responsively more dependable by situational influences, but also to hold that situational influences can make you trust-responsively more durable. The empowering theory defends both the situational enhancement of dependability and the situational reinforcement of durability.

2.2 *The situational enhancement of dependability*

There are three grounds for postulating the likelihood of situational enhancement of dependability. The first is that when I trust you, I display and communicate a belief in your capacity to be motivated by my manifest reliance, thereby encouraging you to prove reliable. The second is that when I trust you to do something, I often make a request, explicit or implicit, that you do it. And the third is that when I trust you I display a good opinion of your dependability, thereby giving you an extra esteem-based motive for not letting me down. In any situation, these effects are liable to make you trust-responsively dependable in a higher degree than would be ensured by the standing capacity of trust-responsiveness that you bring to the situation.

2.2.1 THE ENCOURAGEMENT ARGUMENT

The first argument for the situational enhancement thesis claims that insofar as I trust you to do something, I encourage you to do all you can to prove dependable (McGeer 2008b). This turns on three plausible psychological observations: first, that in trusting you I display and communicate a belief in your capacity to be motivated by my manifest reliance on you; second, that in communicating the belief that you have such a capacity, I may boost your own belief that you have that capacity; and third, that in doing this I may improve the capacity itself, making it more dependable: I may give you the courage to test and thereby develop your own powers of agency.

The first of these observations simply falls out of our account of trust, as distinguished from mere reliance. When I trust you, the manifestness of my reliance means that it must be a matter of common awareness that I am doing so: namely, that I am acting on the belief that you will recognize that I am relying on you and be motivated by that very fact to respond to my reliance. This will be a matter of common awareness insofar as the evidence that supports the claim that I am trusting you is clearly available to each of us, ensuring that each will be aware that I am acting out of trust, aware that each is aware of this, and so on (Lewis 1969). Given that trust requires a common awareness of this fact, I will not only display my belief that you are trust-responsive in the involuntary manner in which my yawning might display my boredom. I will display my belief in an intentional manner, communicating that I believe in your capacity ~~for trust-responsiveness~~ (Grice 1957).

The second observation in support of the encouragement argument turns on the commonsense assumption that the fact that I communicate a belief in a capacity like your trust-responsiveness is likely to increase your own confidence in having that capacity. Short of my being a doting parent, or something of that kind, I credibly testify to signs of your having that capacity when I communicate my belief in your

responsiveness, and stake my welfare on its being true. In effect I say: ‘Have no doubts; you possess the capacity to live up to my trust, to prove reliable in the ways that I anticipate!’ And in saying this—indeed, in staking my own welfare on its being true—I provide good ground for your being more confident of having the responsiveness that I ascribe to you. After all, the signs to which I testify may not be accessible to you yourself; your diffidence or fear may lead you not to credit them; or your experience (or inexperience) may lead you to underestimate what they imply about your own powers of agency.

The third observation may seem more problematic, but it is also pretty plausible and follows quite naturally from the last. If you gain in confidence about having a certain sort of capacity—in particular, the capacity to be moved to action by a consideration like my relying on you to do something—then that is likely to improve the capacity itself, increasing the degree to which you dependably exercise it. It is likely to increase the range in which you are trust-responsive beyond the range ensured by your standing responsiveness to reasons of trust: what we earlier called ‘trustworthiness’. This observation is grounded in the fact, borne out in everyday experience, that becoming more confident about being able to do something X in suitable circumstances—say, about being able to prove reliable in response to someone’s manifest reliance on you—means improving that very capacity; it means that you will do X more dependably across relevant circumstances.⁹

Thus, suppose you have a low opinion of your own moral quality, believing that you are weak and give in to self-interest more than you ought to do; believing in particular that you are less responsive than you should be to others who manifestly rely on you not to let them down. Your self-confidence in your own capacity to do as you ought is likely to improve as others display faith in your capacity to prove trust-responsive. And that in turn is likely to encourage you in your display of such responsiveness.

2.2.2 THE REQUEST-BASED ARGUMENT

The second, request-based argument for the situational enhancement thesis turns on the intuition that when I ask you to do something, thereby presupposing that you will be motivated to comply by my manifest reliance on you—that is, that you will respond in an active or back-up way to the reasons of trust relevant in our relationship—the fact that I am making a request may have a force over and beyond the persuasive power of the reasons themselves. Moving away from trust for a moment, suppose I believe that you are fully aware of all the reasons there may be for helping me in some manner. Even in the presence of such a belief it may make sense to ask for the help I need. Doing so will amount to recognizing the burden that helping me is likely to

⁹ ~~Here as elsewhere there is some interpretive latitude about how to understand the disposition and about where to draw the line between dependability and durability. Thus the text assumes that having the capacity to do X at any level of dependability is responsive to your confidence about being able to do X. If it were unresponsive, then the absence of confidence, assuming it affects your performance, would count as a disrupter of the capacity, affecting the level of its durability rather than its dependability.~~

impose, acknowledging that you have no obligation to bear that burden, and asking your permission nevertheless to impose it. If you do not then decline the burden, but implicitly accept it, you will effectively have promised to provide the help I require. And so you will have an additional normative reason, and presumably an additional motive, for actually providing that help.

When I manifestly rely on you to act in a certain way, I may often make an explicit request to do what I am relying on you to do, as when I ask you to keep an eye on my computer. And even where I do not make an explicit request, reliance may be so manifest as to constitute an implicit request. By analogy with the help case, that means that in either of these cases I effectively ask your permission to impose on you the burden of proving reliable, acknowledging that it will indeed be a burden and that you have no obligation to bear it. And if you do not decline to bear the burden—if in effect you accept it—then that means that there is further reason to bear it: that otherwise you will breach an implicit promise. However morally binding you consider that reason, it is very likely that you will be motivated by it; no one wants to be cast as unmoved by their promises. And the motivational effect of the promise will be to make you more responsive to the reasons or requirements of trust.¹⁰

2.2.3 THE ESTEEM-BASED ARGUMENT

We have looked at two arguments for believing in the situational enhancement thesis, according to which I may increase the dependability of your trust-responsive disposition by the very fact of displaying trust in you. Both of those arguments suggest that I may increase your dependability by introducing factors that we can each acknowledge openly (Baier 1986; Jones 2004). I do not undermine my impact by making clear that I am hoping to have an impact, whether in expressing my faith in the capacity that I take you to have or in making a more or less explicit request to you. The factor invoked in each case is publicly avowable between us. Consistently with operating as described, it may be something of which we are each aware, aware that each is aware, and so on in the usual regress of common awareness.

The factor introduced in the third argument for situational enhancement is different in this regard from the other two (McGeer 2008b). It consists in the fact that when I display trust in you to act in a certain manner I give evidence, by regular accepted criteria of attitude ascription, of having a good opinion of you. And I thereby provide you with a new motive for behaving as I am manifestly relying on you to behave; that by doing so, you will reinforce the good opinion I exhibit and secure an intuitively attractive reward in the economy of esteem (Brennan and Pettit 2004). This reward will be all the

¹⁰ A distinct argument might appeal to the thought that second-person claims constitute a unique source of moral demand, distinct from any impersonal demand to improve the world or indeed any personal demand to act in a morally righteous way (Darwall 2006). If this is correct, then the request-based argument need only assume that in asking for your trust-responsiveness, I create a second-personal reason why you should prove reliable.

greater of course, and the new motive will be all the more powerful, if I trust in you as a public act; in this case, you will be able to secure the good opinion of witnesses, and not my good opinion alone, by proving reliable.

This argument suggests that there is a certain cunning in trust (Pettit 1995). Our accepted criteria of attitude ascription support the view that when I trust you to act in a certain manner, that must be because I hold a good opinion of you. Those criteria embody ‘the fundamental attribution bias’ (Jones 1990): the tendency to explain responses like my trusting you in terms of some manifest, standing disposition such as the belief that you are trustworthy. As a matter of fact, however, I may not be acting on such a belief in trusting you. Rather I may be grounding my trust in the belief, first, that you will care for the effects on your reputation with me and with any witnesses of your proving to be reliable; and, second, that this concern will almost certainly lead you to act as I am purportedly trusting you to act.

There is more to be said about the dynamics of esteem in supporting trust and, indeed, other such overtures between people. But these remarks should be sufficient to identify a third factor that supports the idea that by trusting you I may do more than activate a pre-existing disposition of trustworthiness to be moved by the manifest reliance of others; I may induce a higher level of dependability in your trust-responsiveness. My display of trust may not only constitute a motivating request for help and an intentional form of encouragement to provide that help. It may also link your provision of help—your acting as I rely on you to act—with an attractive prospect of securing esteem and standing in my eyes and in the eyes of witnesses.

2.2.4 A QUESTION

We describe the three factors surveyed here as enhancing the dependability of your standing disposition to respond to reasons of trust: that is, as making it more dependable. But why not take them to complement your standing disposition by putting independent forces in place that have a parallel, if convergent effect? Why not take them to do for your trustworthiness what your self-interest does for your honesty, when it is clearly in your interest to tell the truth?

The answer is that the three factors in our examples each serve to make you more responsive to reasons of trust as such, not just to make you more likely to act in the way that, as it happens, reasons of trust support. The encouragement I provide in trusting you boosts your existing capacity, whatever that involves, to be motivated by reasons of trust. The request I make in doing this provides you with a further, presumptively motivating reason—that it will enable you to keep an implicit promise—to act in whatever manner those reasons require. And equally the prospect of esteem that I put in place when I trust you provides you with a further, motivating reason to be sensitive to the reasons of trust as such: after all, you can only expect to win credit from me by proving yourself trust-responsive in this sense.

2.3 Modelling situational enhancement

The upshot of this picture is that your responsiveness to evidence of trust in a given instance, your capacity to be moved to action by my manifestly relying on you, may be a function of two factors.¹¹ First, the standing sensitivity to reasons of trust that you bring to that encounter: that is, your trustworthiness. And, second, the situational sensitivity to my presence that enhances your trustworthiness, increasing the dependability of your disposition to act as the reasons of trust require. We might picture the play of these factors as shown in Figure 2.1.

If your trust-responsive dependability is a function of two forces, of course, then it becomes possible that in a given case it may result from different combinations of those forces. The two sensitivities may combine in different measures to produce responsiveness and any degree of responsiveness may be realized via any of a range of equivalent combinations. In some combinations sensitivity to reasons of trust will be high and sensitivity to the presence of a trustor low, in others the reverse will hold, and in still others the factors will be more or less equal.

Thus we might represent your trust-responsive dependability in a given instance on a graph in which the vertical axis depicts sensitivity to reasons of trust, the horizontal sensitivity to the presence of a trustor (see Figure 2.2). Take a given combination as a point on the graph and now connect up those points that yield the same degree of trust-responsiveness. Such points will connect in an equivalence curve—one of the diagonal lines in the diagram—in the space of trust-responsive dependability. We may characterize any of a range of curves, each corresponding to a different degree of dependability: the lines to the upper right in the diagram will represent higher degrees, those to the lower left lower degrees, in analogy with the familiar picture of indifference curves in the space of utility. One of those curves will presumably characterize your degree of trust-responsiveness in the case in question. And on the curve that identifies you, there will be a point that represents your particular combination of sensitivities in that case.




Figure 2.1 Responsiveness to evidence of trust

¹¹ We borrow in discussion here from McGeer and Pettit (2015).




Figure 2.2 Trust-responsive dependability

2.4 *The situational reinforcement of durability*

If the considerations rehearsed so far are sound, they suggest that situational influences are likely to make you trust-responsively more dependable than could be explained by the standing capacity of trust-responsiveness that you bring to the situation. But the empowering theory of trust, as we understand it, holds also by a second thesis: that situational influences are likely, not only to make your disposition to respond to reasons of trust more dependable, but also to make it more durable; they are likely, not just to enhance the disposition, but to reinforce it. The idea is that just as situational influences can increase the range of scenarios over which that disposition is dependably triggered, so they can increase the range of scenarios in which it is durable enough to survive potential disrupters.

The easiest way to support the reinforcement thesis is to consider the case where you have developed a more or less dependable, standing disposition of trust-responsiveness, perhaps under the long-term influence of situational sensitization. In our terms, you have become a more or less trustworthy person, not just someone who is trust-responsive. The considerations rehearsed in support of the enhancement thesis argue that situational influences can also reinforce any standing disposition of trust-responsiveness that already exists.

One way in which those influences may reinforce your trust-responsiveness, of course, is by habitually reminding you of the lessons they teach: that if you do not prove reliable, you must often count as breaking a promise to a trustor, you must disappoint the faith in you that they display in encouraging you, and you must expect to lose the good opinion and reputation that they give you the chance to earn. The causal effect of being constantly under such tutelage will surely be to maintain and reinforce in you the disposition to respond to reasons of trust. If, for example, your standing trust-responsiveness has begun to decline, say as a result of losing faith in yourself, an exposure to such influences from others—in particular, perhaps, their influence in encouraging you—ought to help to restore the durability of that disposition. This exposure will help to maintain your standing trust-responsiveness in the way in which booster shots of a vaccine may help to maintain your active immunity to a disease.

But apart from the causal impact of situational influences in maintaining or reinforcing your trust-responsiveness, there is also a further consideration that argues for

the reinforcement thesis. Suppose that your standing disposition to respond to reasons of trust is so dependable that there is no room for further enhancement by situational influences. Even in that case, there is a role that those influences can continue to play in making the disposition more durable.

Once again, the ring of Gyges story brings out that role nicely. Suppose, as seems reasonable in many cases, that you would cease to be trustworthy—you would lose any standing capacity to respond to reasons of trust—if you enjoyed the total undetectability and impunity that the ring would confer; those factors would disrupt your trustworthiness, suspending it or even obliterating it. What does that teach us about the factors that play a role in keeping the disposition in place in the ordinary world? Presumably, that your exposure to others is essential to your continuing to have the disposition, to its remaining a durable part of your psychology, and plausibly, that your exposure to the sorts of situational influences listed has a major part in securing that result.

How are we to conceive of the role that those factors play, assuming that you are already fully trustworthy? You may not be focused on the goal of proving yourself a promise-keeper, living up to the encouragement of others, or winning esteem in their eyes. And you may not even be aware of such influences, let alone attentive to them. All that need be true is that the influences have a virtual impact on your thinking and behaviour. They are there to put you back on course—there to maintain and reinforce your trust-responsiveness—should you fall away. The ordinary world is one where any tendency to give into temptation, or any step in that direction, is likely to trigger those influences and restore the disposition in you. And that shows up in the fact, as we presume it to be, that in the Gyges world there is nothing to block any such fall from grace and much indeed to prompt a fall: namely, the salient fact that you can do as you wish with total impunity.

On this picture, the situational influences may not have any causal effect in ensuring that your trustworthiness remains durably present. What they have rather is a virtual effect. They may play no causal role in keeping your trustworthiness in place—it survives as a matter of deliberative habit—but if they were absent then, sooner or later, that standing responsiveness to reasons of trust would fall away. Their absence would make self-seeking deliberation a salient option and would almost certainly undermine the disposition.

This picture implies that the situational influences to which you are subject control for the durable presence of the disposition. Think of how the cowboy controls his cattle in the classic western movie. He lets them follow their head—he does nothing causally to direct them—but is ready to intervene should any head of cattle wander off track. The situational influences that surround you in your role as a potential trustee have the same sort of non-causal control over your trust-responsiveness. They give that disposition a durable presence that it would not enjoy in their absence. And they do this just by being there, ready to play a causal role if they are needed to maintain the disposition.

3. Trust-Responsiveness on the Empowering Theory

3.1 *The evocative ascription of trust-responsiveness*

When I trust you I display, and indeed communicate, the belief that you are trust-responsive, as we saw in section two. But that means that trusting you has the force of an ascription of trust-responsiveness. It amounts to saying in an ascriptive mode: ‘you can respond to the fact that I am relying on you; you have the capacity to be moved by that consideration; you can prove yourself reliable.’

On the empowering view of trust, however, this ascription of capacity has a very unusual character. That you have the capacity of trust-responsiveness in a given situation means that you are disposed, however durably, to respond more or less dependably to reasons of trust. Whether or not you actually fail to act on the capacity, it must be true that in a range of variations on actual circumstances—in particular, ones where reasons of trust remain in place, where no outweighing considerations are introduced, and where no novel perturbation affects you—you would respond as appropriate. But if the empowering view is sound, then when I ascribe such a capacity to you in displaying trust, I am likely to have an effect on the very capacity I ascribe.

The likelihood of having such an effect is surely going to be salient in the culture.¹² And so I cannot ascribe the capacity in a purely descriptive spirit, as if I were recording a pre-existing state of affairs. In asking you to take down a bottle from a shelf I cannot reach, I am implicitly ascribing a height to you that I do not possess myself and in doing so I am recording what I take to be a pre-existing state of affairs and making the ascription in a purely descriptive spirit. The ascription of trust-responsiveness that I implicitly make in trusting you is clearly not like that. It is akin to the ascription of a capacity—say, a capacity to reach the top of a hill on your bicycle—when I call out in encouragement: ‘You can do it’. The ascription is inevitably designed to elicit the very capacity that it attributes.

Descriptive reports are often contrasted with performative reports. In the case of a performative, the utterance brings about the very state of affairs on which it reports, as in examples like ‘I resign my post’ or ‘I beg your favor’ (Lewis 1983, ch. 12). But as the ascription of a capacity that I implicitly make in trusting you is unlike a descriptive report, so it is also unlike a performative utterance: I do not make you responsive to evidence of trust just by ascribing such responsiveness to you.

The descriptive and performative alternatives, however, do not exhaust the options. There is a third possibility that our discussion makes salient. This is that the ascription of trust-responsiveness implicit in my displaying trust in you is an *evocative* ascription, as we may call it: an ascription that is designed to help bring about the very state of affairs on which it reports. The idea is that in communicating the message that you are

¹² This may not mean that we regularly pay attention to the effect but only that the frequency with which we resorted to ascribing trust-responsiveness would be affected if it became clear that the effect did not generally materialize.

trust-responsive and can prove reliable, I do so with complex intent. I mean to try to elicit your capacity to prove reliable, or at least to support it, in the very act of recording its presence. Whether or not I make this explicit, I am effectively exhorting you to prove reliable. I am speaking in a way that I may reasonably expect to have a direct effect in enhancing the dependability of your trust-responsiveness—and an indirect effect perhaps in reinforcing its durability.

How are we to think of a capacity that has such a plastic form that I can ascribe it to you in this evocative fashion? The question is worth exploring as it turns out that there are three significant lessons that the argument teaches us about how to think about the capacity. The first is that the capacity is context dependent; the second that it is a work in progress; and the third that assigning the capacity is subject to moral as well as epistemic norms.

3.2 Trust-responsiveness is context dependent

The first clear implication of the empowering theory is that your capacity to prove reliable in dealing with a trustor like me is not fixed solely on the basis of how you are in yourself. It may be fixed, not just by your standing sensitivity to the reasons of trust—that is, by your pre-established trustworthiness—but also by your situational sensitivity to me, your trustor. And even if it is wholly fixed by your standing sensitivity to reasons of trust, that sensitivity is itself kept in place by your situational sensitivity to me, and to other trustors. The capacity of trust-responsiveness, on this picture, is the product or resultant both of how you are in yourself and of how things are in the world you occupy.

It is common to assume that any disposition you possess—and so any capacity linked with that disposition—must be grounded in categorical facts; it cannot be a bare disposition, as it is sometimes put (Dummett 1973). The basis of the assumption is that it is hard to imagine that two otherwise identical subjects should be the bearers of different dispositions. What could explain the difference in disposition other than a difference of a more categorical kind?

But when it comes to dispositions like trust-responsiveness, and human capacities more generally, this common assumption takes a particular form. It is taken to imply that as between two subjects who differ in such a disposition, the categorical difference that explains the divergence must be a difference in their inherent make-up. If our observations about trust-responsiveness are correct, however, this is not so. How trust-responsive you are in dealing with me may be a function of many factors: my nature, our relationship, and the company that we share with others, such as others who serve as witnesses to my relying on you in a certain way. In short, it may be a function of the context in which you operate as well as a function of your inherent character. It may be a context-dependent or ecological capacity, not a capacity that is fixed only by how you are in yourself.¹³

¹³ This notion of an ecological capacity resembles one defended in another context by Manuel Vargas (2013); we borrow the term ‘ecological’ from his book. See also McGeer and Pettit (2015). The fact that the

3.3 *Trust-responsiveness is a work in progress*

The standard view of trust reifies trust-responsiveness and suggests that its presence or absence is an on or off matter, so that there is a clear dividing line between those who are worthy of trust and those who are not. The capacity is equated with a disposition to appreciate and respond to reasons of trust and like any such disposition it may come in various degrees of dependability and durability; it may dispose the agent more or less durably to respond to reasons of trust more or less dependably. But the idea in the standard view is that there is a certain threshold of durability and dependability at which this capacity materializes and that every agent reaches this or fails to reach it; there is no middle ground. It may take educational and moral development for an agent to pass the threshold. But given that development, the capacity is fixed in place, with a suitable degree of dependability and durability, and the agent joins the ranks of the trustworthy.

There may well be a certain threshold of responsiveness to reasons of trust, both in durability and dependability, which we require those who count as trustworthy to pass. But that should not lead us to adopt the standard view. For on our picture, the responsiveness may enjoy that degree of dependability under equivalent mixtures of two factors: one, a standing sensitivity to reasons of trust as such; and two, a situational sensitivity to the presence of a trustor. And on our picture, the responsiveness may be more or less durable, depending on the effectiveness of situational influences on the trustee.

Once we recognize that trust-responsiveness has this ecological character, it becomes difficult to think of the capacity involved as something that is there or not there in an agent, depending on the threshold reached. By our account, the responsiveness or capacity required may materialize at a given level of dependability on the basis of any mix of standing and situational sensitivity, and it may materialize with one or another degree of durability. It would make sense to posit the sort of threshold envisaged in the standard theory if trust-responsiveness were something established in the make-up of a person. But this scarcely makes sense on the basis of the assumptions defended here.

The image suggested by our approach is that the capacity that makes you fit to be trusted is as likely to be a work in progress—a work in which others play a role and not just you alone—as it is to be a finished achievement: something that is fixed in place, once for all. The moral life promises to be agonistic in nature, involving a continuous struggle to maintain an appreciation of the reasons relevant in trust and to sustain your responsiveness to those reasons (McGeer 2008a). It is likely to require you to call on all the reserves at your disposal, including reserves that only suitable social interactions can provide, in order to maintain your status as someone it makes sense to trust.

capacity is context-dependent might lead us to treat ascriptions as indexical, so that what I ascribe is always responsiveness-to-me, but it would also be consistent with treating them as ascriptions of responsiveness, now to this sort of audience, now to that. See Hawthorne and Pettit (1996).

There may be an attraction, of course, in the prospect of achieving an increasing level of sensitivity to reasons of trust and a decreasing dependence on sensitivity to the presence of one or another trustor. This is the appeal of becoming a virtuous person—specifically, a trustworthy person—in the Aristotelian sense of virtue: someone with a potential to act appropriately that is not hostage to the fortune of circumstance; someone who not only enjoys excellence or *areteia* but enjoys it with a certain self-sufficiency.

But, however appealing in other ways, the Aristotelian ideal may be a myth in at least one respect. Even if you have a standing disposition to respond dependably to reasons of trust—even if the dependability of your trust-responsiveness is grounded in your nature—you may depend on context, and not just on your inherent nature, for the durable presence of that disposition. This is a lesson that the ring of Gyges teaches. Or a lesson that it teaches, at any rate, on the assumption that many of us would not maintain even a hard-won degree of trustworthiness—a standing capacity to respond to reasons of trust—if we could manage with impunity to break promises, flout the encouragement of others, and perform manifestly shameful acts. Even if some of us have achieved the context-independent dependability associated with the Aristotelian ideal, so the message goes, we are unlikely to have achieved a corresponding degree of context-independent durability.

3.4 *Ascribing trust-responsiveness is subject to moral norms*

Since the ascription of a capacity like trust-responsiveness is a judgemental exercise, bound to the requirements of truth, it is certainly subject to epistemic or cognitive norms. I ought not to ascribe a capacity to appreciate and act on the manifest reliance of others in the absence of all evidence that you have a suitable disposition to respond to reasons of trust. Certainly, for example, I ought not to ascribe such a capacity if experience shows that you routinely let others down, even when they go out of their way to make it clear that they are relying on you.¹⁴

But to ascribe trust-responsiveness, by the argument of this paper, is often to enhance or reinforce the capacity, helping to make it more dependable or durable. And that means that it could not be right to let the ascription of trust-responsiveness be governed by epistemic or cognitive norms alone. That would be appropriate only if the ascription amounted to nothing more than the report of an unobtrusive observer. But in interacting with you as trustor to trustee—or would-be trustor to hoped-for trustee—I am never an observer of that merely spectatorial kind.

The lesson of this observation is that while I should certainly not breach cognitive norms in crediting you with trust-responsiveness, whether by explicit ascription or by implied assumption, I should be guided in the slack that those norms leave by norms of a moral or ethical kind. I should let myself be led to attribute this capacity to you on the

¹⁴ Of course even in this case I might have a developmental rationale for manifesting reliance, and at least purporting to ascribe trust-responsiveness to you: I might plan to nudge you over time towards being trust-responsive; see the discussion at the end of section one.

basis of how far this is likely to be productive, helping to enhance or reinforce the capacity I posit. I should take into account the extent to which you are likely to be scaffolded by my attribution and to perform in a manner that is immediately beneficial to me and that may be more generally beneficial in reinforcing your trust-responsiveness.

Ascribing trust-responsiveness to you, at least in the way in which this is encoded in a display of trust, means investing in you, as it is said, holding out the hope that you will live up to the expectations associated with that ascription (McGeer 2004; Pettit 2004). When I invest this hope in you, I have reason to think that the relationship I thereby announce will motivate you to live up to my expectations, proving worthy of my faith. But that means that there are moral norms that should guide me in displaying trust in you and ascribing a suitably dependable and durable degree of trust-responsiveness. The ascription of such responsiveness may be subject to constraints of evidence and a fit topic for epistemology. But it is also subject to moral standards and is a suitable topic for ethics.

We end with an illustration of this last lesson (McGeer 2008b). In George Eliot's novel *Middlemarch*, the local doctor, Tertius Lydgate, has fallen on bad days, guilty by association with the villainous Nicholas Bulstrode and is abandoned in disgrace, left to despair, by the suspicious and gossip-mongering townsfolk. The high schemes he had entertained of ushering in new medical discoveries and practices are certainly in tatters, and his prospects otherwise look dim. But at this crucial juncture Eliot's heroine, Dorothea, calls his coterie of friends to account, reminding them that true friends have a duty to support one another, so that Lydgate might rise to the challenges he faces (Baker 1987). Admonishing them, she declares:

'And would you not like to be the one person who believed in that man's innocence, if the rest of the world belied him? . . . I should not be afraid of asking Mr Lydgate to tell me the truth, that I might help him. Why should I be afraid? I . . . [could] ask for his confidence; and he would be able to tell me things that might make all the circumstances clear. Then we would all stand by him and bring him out of his trouble. People glorify all sorts of bravery except the bravery they might show on behalf of their nearest neighbours.'

Due to her intentions, Dorothea seeks Lydgate out to discover the full story behind his disgrace, and while his tale is far from completely self-exonerating, he is able to confess to, and take responsibility for, his own follies and weaknesses thanks to Dorothea's continuing display of trust in him. Better, he is able to leave her with a renewed faith in his own capacity to overcome his difficulties, thereby living up to the faith she expresses in him. As Lydgate himself remarks, 'you have made a great difference to my courage by believing in me. Everything seems more bearable since I have talked to you' (Eliot 1994: 768). These empowering effects of Dorothea's trust make perfect sense in light of the themes we have been charting in this paper.¹⁵

¹⁵ We were helped enormously in revising this paper through the discussion of participants and a commentary provided by Wlodek Rabinowicz at a moral philosophy conference in Beaune, France in June 2015; by delegate input at the Kioloa Moral Philosophy Workshop in Jan. 2016; and by copious comments from Paul Faulkner and Thomas Simpson.

References

- Alfano, M. (2014). What Are the Bearers of Virtues? In *Advances in Experimental Moral Psychology*, ed. H. Sarkissian and J. C. Wright. London, Bloomsbury Academic: 36–54.
- Baier, A. (1986). Trust & Antitrust. *Ethics* 96: 231–60.
- Baker, J. (1987). Trust and Rationality. *Pacific Philosophical Quarterly* 68.
- Brennan, G. and P. Pettit (2004). *The Economy of Esteem: An Essay on Civil and Political Society*. Oxford, Oxford University Press.
- Darwall, S. (2006). *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge, MA, Harvard University Press.
- Dummett, M. (1973). *Frege: Philosophy of Language*. London, Duckworth.
- Eliot, G. (1994). *Middlemarch*. New York, Penguin Books.
- Grice, P. (1957). Meaning. *Philosophical Review* 66: 377–88.
- Hawthorne, J. and P. Pettit (1996). Strategies for Free Will Compatibilism. *Analysis* 56: 191–201.
- Holton, R. (1994). Deciding to Trust, Coming to Believe. *Australasian Journal of Philosophy* 72: 63–76.
- Jones, E. E. (1990). *Interpersonal Perception*. New York, Freeman.
- Jones, K. (1996). Trust as an Affective Attitude. *Ethics* 107: 4–25.
- Jones, K. (2004). Trust and Terror. In *Moral Psychology: Feminist Ethics and Social Theory*, ed. P. DesAutels and M. U. Walker. Lanham, MA, Rowman and Littlefield: 3–17.
- Kaplan, D. (1978). On the Logic of Demonstratives. *Journal of Philosophical Logic* 8: 81–98.
- Lewis, D. (1969). *Convention*. Cambridge, MA, Harvard University Press.
- Lewis, D. (1983). *Philosophical Papers: Volume I*. Oxford, Oxford University Press.
- McGeer, V. (2004). The Art of Good Hope. *Annals of the American Academy of Political and Social Science* 592: 100–27.
- McGeer, V. (2008a). The Moral Development of First-Person Authority. *European Journal of Philosophy* 16: 81–108.
- McGeer, V. (2008b). Trust, Hope and Empowerment. *Australasian Journal of Philosophy* 86: 237–54.
- McGeer, V. and P. Pettit (2015). The Hard Problem of Responsibility. In *Oxford Studies in Agency and Responsibility*, ed. D. Shoemaker. Oxford, Oxford University Press. Vol. 3: 160–88.
- Olsaretti, S. (2004). *Liberty, Desert and the Market*. Cambridge, Cambridge University Press.
- Pettit, P. (1995). The Cunning of Trust. *Philosophy and Public Affairs* 24: 202–25; reprinted in P. Pettit (2002). *Rules, Reasons, and Norms*. Oxford, Oxford University Press.
- Pettit, P. (2004). Hope and its Place in Mind. *Annals of the American Academy of Political and Social Science* 592: 152–65.
- Pettit, P. (2015). *The Robust Demands of the Good: Ethics with Attachment, Virtue, and Respect*. Oxford, Oxford University Press.
- Smith, M. (2003). Rational Capacities, or: How to Distinguish Recklessness, Weakness and Compulsion. In *Weakness of Will and Practical Irrationality*, ed. S. Stroud and C. Tappolet. Oxford, Oxford University Press: 17–38.
- Vargas, M. (2013). *Building Better Beings: A Theory of Moral Responsibility*. Oxford, Oxford University Press.