



This is a repository copy of *Finding trust in government*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/140229/>

Version: Accepted Version

Article:

Faulkner, P.R. (2018) Finding trust in government. *Journal of Social Philosophy*, 49 (4). pp. 626-644. ISSN 0047-2786

<https://doi.org/10.1111/josp.12262>

This is the peer reviewed version of the following article: Faulkner, P. (2018), Finding Trust in Government. *J Soc Philos.*, which has been published in final form at <https://doi.org/10.1111/josp.12262>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Finding Trust in Government

Much of what is good, and civic, in society rests on trust. Trust, it has been claimed, furthers economic prosperity and growth (Putnam 1993 and Fukuyama 1995); fosters good government (Knack 2002); and lowers level of corruption (Uslaner 2013). Moreover, some basic level of trust seems to be necessary for society to function. For consider what would remain if trust were removed and replaced by distrust. Were this so, we could not enter any form of exchange, for why risk handing over goods without payment, or payment without goods? And we would even seem to lack reason for thinking others won't attack us, given our weakness and the obvious utility our goods have for others. It might sound strained to say we trust others not to be "unpromptedly aggressive", but this is because we live in "better times" (Williams 2002, 89), where trust, in this respect, can be taken for granted. Remove this taken-for-granted background, as is being supposed, and what we imagine is a state of "warre", "of everyman against every man" (Hobbes 1651, 185).

So trust is necessary for society to function and underwrites much of what is good in society, but working out the role of trust in our social and political lives is no straightforward matter. In part, this is because empirical work does not proceed from a theory of trust. Thus, in a

review of empirical work on trust, Nannestad (2008, 415) observes that this work doesn't tend to proceed from a clear idea of what trust is; "[r]ather, trust is taken to be what is measured by one or more survey questions, or by subject's observable behaviour in certain experimental games".

And in part, this is because different theories of trust imply different accounts of the social and political role of trust. Thus, Hardin (2002, 152) takes trust to be belief about what is in the trusted's interest, and concludes that "trust in government is conceptually and epistemologically impossible for most citizens in large modern societies". While Uslaner (2002, 7) takes trust to be an essentially optimistic attitude and concludes that trust is primarily of social, rather than political, importance since "[t]rust in people and trust in government have different roots" (see also (Lenard 2005, 370)).

In fact, trust can be a belief about interest *and* an optimistic attitude, and the starting point of this paper is a philosophical theory of trust that makes sense of this conjunction. It does so by differentiating between different attitudes of trust and different objects of trust. In the political domain, a further distinction then needs to be added. Trust can be vertical: directed towards government and government representatives. And it can be horizontal: directed towards other citizens. Only when both forms of trust are present, Lenard (2015, 314) argues, can a state be "characterized by a climate of trust". Through applying these distinctions, deriving from a theory of trust, the present paper aims to say something about the role of trust in the political domain.

What results shows Hardin to be wrong: we can trust government, even if we cannot calculate a

government representative's interests. And Uslaner to be wrong: trust is essential to democratic government, even if it plays a fundamental social role.

1. A Theory of Trust

First, the theory of trust.¹ Trusting is both something we do and an attitude that we can have and take. The *act of trusting* is one of putting oneself in a position of depending on something happening or someone doing something. Thus, and for instance, my delaying my departure is also the act of trusting my car to start; and my waiting in this café is also the act of trusting you to turn up as agreed. I am in a position of dependence in these cases because if my car doesn't start or you don't show, I will be troubled in practical and other ways. An act of trusting is thereby one of relying: in trusting we rely on something happening or someone doing something. But, and crucially, trust differs from reliance in that it is necessarily willingly; although reliance can be willing – one can be confident that someone or thing will prove reliable — it can also be forced: one might have no choice but to take the least bad option. This is not the case with trust, which, as Williams (2002, 88) observes, “involves the willingness of one party to rely on another to act in certain ways”. The *attitude of trust* is then characterized as an attitude towards this dependence, or

¹ The theory presented is developed in Faulkner (2007b), Faulkner (2007a), Faulkner (2010), Faulkner (2011), Faulkner (2012), Faulkner (2014), Faulkner (2015), and Faulkner (2017)

reliance, that *explains its willingness*. Thus, in the political domain, a climate of trust ensures that there will be “a general willingness among citizens to comply more or less voluntarily with democratically determined legislation” (Lenard 2015, 353)

How trust explains willing reliance is then determined by its nature as an attitude and here two distinctions need to be drawn: the first is metaphysical and concerns its object; and the second psychological and concerns its kind.

First, *the objects of trust*. The attitude of trust can be an attitude towards *either* some particular instance of dependence or depending on some particular person; *or* the prospect of such dependence more generally. To elaborate, the object of trust can be some particular dependence; we speak of trusting our car to start, trusting you to turn up on time, trusting the ice to hold our weight, trusting another to deliver the goods as promised, and so on. Focussed on a particular instance of reliance, trust is a three-place predicate with the form ‘X trusts Y to ϕ ’. Here ‘X’ refers to the trusting party, ‘Y’ the trusted person or thing, and ‘ ϕ ’ the action or event that X relies on happening. Three-place trust is a metaphysical hybrid: it combines an event or action with an attitude towards it. In this respect, trust might be compared to knowledge, which on the standard account is equally metaphysically hybrid in combining an attitude with the fact that makes it true. But we also speak of trusting particular persons: it is not that I merely trust you to turn up on time

or deliver the goods as promised, rather I trust you in these respects simply because I trust you more generally.

When the focus is shifted from a particular instance of reliance to reliance on a particular person, trust becomes a two-place predicate with the form 'X trusts Y'. In this case, 'X' again refers to the trusting party but 'Y' can only refer to a trusted person. One would not say 'I trust my car' or 'I trust the ice', or if one did say this, it would be short-hand for some three-place claim such as 'I trust my car to start' or 'I trust the ice to hold my weight'. (To understand the restriction of two-place trust to persons, one needs to consider the second distinction between the two kinds of psychological attitude trust can be, of which more shortly.) However, the two place 'X trusts Y', where 'Y' is a person, cannot be similarly translated into a three-place claim. In trusting another person one demonstrates a general willingness to rely on them, but there can be no precisification of this since one would never trust someone in all ways and there is no specific way that one must be willing to rely on them for it to be true that one has this general willingness. While three-place trust is metaphysically hybrid, two-place trust is not: it is simply an attitude that one can have towards persons.

However, it is equally possible to remove this reference to a particular person and give trust a one-place grammatical form: 'X trusts' — or maybe the grammatically better 'X is trusting'.

“[W]e do”, Uslaner (2002, 22) observes, “speak of trusting people generally”. And this form does not seem reducible to two place trust for the same reasons that two-place trust is not reducible to three-place trust: it will not be that X trusts everyone, and there is no specific set of people that X must trust for it to be true that X is trusting. Rather, ‘X is trusting’ seems to make a different claim: that X has faith in people, in some “generalized other”, as Uslaner (2002, 24) says. But while ‘X is trusting’ and ‘X trusts Y’ seem to be different and unique statements in the same way that ‘X trusts Y’ and ‘X trusts Y to ϕ ’ seem to be different and unique statements, there is some implication in both cases from the former to the latter but not the other way round. So of the three predicates, the one-place predicate is arguably the most fundamental. Thus, the heart of our notion of trust seems to be simply an attitude of trust, which may, but need not, take specific persons as its object, and which can support, but need not, the act of relying on persons. So this distinction between the objects of trust identifies, what Uslaner would call, two *particularized* forms of trust — the three-place ‘X trust Y to ϕ ’ and the two-place ‘X trust Y’ — and one *generalized* form of trust — the one-place ‘X trusts’ or ‘X is trusting’.

Second, the psychological kinds of trust. The simplest attitude that could explain a willingness to rely would be belief. If X believed that Y would ϕ , then X’s relying on Y ϕ -ing would be unproblematic. And the attitude of trust can be no more than this: a positive belief about outcome. In so far as it is no more than belief, both persons and things can be trusted. It is because

I believe my car will start, that I am willing to rely on its doing so; and to say that I trust it will start is just to say I have this belief in a situation where I rely on its starting. Our interactions with persons can be no different. We can and do make predictions about how people will behave and then rely on them in the light of these. Cases where we do so, we might describe as trusting. In so doing we identify a very thin attitude of trust; ‘thin’ in that there is nothing evaluative about the correlative judgement of trustworthiness. To say that a person is trustworthy in the associated sense is just to say that they are predictable and so can be relied on in the way one does rely on them. The attitude of trust thus instantiated is simply the belief that things will work out in this case. This conception of trust, Braithwaite (1998, 344) remarks is the “dominant social science conception”, and it is what Uslaner calls *strategic trust*.² Hardin’s conception of trust has this strategic form: he proposes that we trust when reliance is supported by a particular belief. X trusts Y to ϕ when X believes that Y will ϕ , and believes this because he judges that ϕ -ing in Y’s interest. Since Y’s ϕ -ing is also in X’s interest, trust is thereby a matter of *encapsulated interest*: “I trust you because I think it is in your interest to attend to my interests in the relevant matter” (Hardin 2002, 2).

² It is dominant, I suppose, because conceived in this way trust can be theorized using rational choice theory. See for instance (Kreps, Milgrom, Roberts and Wilson 1982), (Glaeser, Laibson, Scheinkman and Soutter 2000) and most of the papers in (Gambetta 1988)

This is fine as far as it goes. The problem two-fold. First, trust as encapsulated interest, or as strategic more widely, can only lead to cooperation given an on-going exchange or with a known interaction partner, where the reliant party X has some specific grounds for thinking that Y will ϕ . It won't allow for cooperation more generally. And it is limitations on when we can judge what lies in others' interests that makes trust in government, according to Hardin (2002, 152), "conceptually and epistemologically impossible". Second, it strains the notion of interest to claim that the grounds that X might have for relying on Y ϕ -ing are solely grounds for thinking that it is in Y's interest to ϕ . And this is because belief is not the only attitude that can explain why reliance is willing. A normative expectation — as opposed to a subjective one — could equally do so in certain circumstances. Take my waiting for you in the café. I don't merely predict you will arrive shortly, I also expect this *of you* or think that you *ought* to arrive shortly; after all, given that we made the plan we did, you have a reason for thinking that I am here waiting. Thus, our ordinary talk of trust takes it, in Uslaner (2010, 112) terms, to have "a moral dimension". Of course, the normative expectation that someone ought to do something does not entail the belief that they will do it. We know that on many an occasion people don't do what they ought to do. However, take a specific case where one has an expectation of someone, a case like my expecting your arrival, in such a case we will presume that the expectation will be met *unless* we have particular reasons for thinking otherwise. (You might, for instance, have a history of keeping me waiting, and in this

case I might both think that you ought to turn up shortly and yet believe that you probably won't; but then this would not be case of my trusting you to turn up shortly.)

Thus, when X trusts Y to ϕ , X's attitude can be simply the belief that Y will ϕ — the expectation characteristic of trust can be merely subjective — but X's expectation can also be normative. Trust, in Uslaner's terms can be *moralistic*. (Though this "is a poor terminological decision" (Lenard 2005, 365) since we can trust people to do immoral things: so the normativity is social rather than moral (see (Faulkner 2010)). Substituting a normative expectation for the subjective one gives a thicker notion of trust, where it is 'thicker' in that the associated notion of trustworthiness is partly evaluative. Such a thick concept is probably the dominant conception of trust within philosophy (see, for instance, (Baier 1986), (Jones 1996), (Becker 1996), (Govier 1992), (Holton 1994) and (Hawley 2014)). As I understand it, the content of this normative expectation is that the trusted party will take one's reliance as a reason for doing what one relies on them doing (see also (Jones 1996), (McGeer and Pettit 2017) and (Darwall 2017)). That is, when X trusts Y to ϕ , X's attitude can be the normative expectation placed on Y that Y will ϕ *because* X relies on Y ϕ -ing. And insofar as X has no reasons for thinking Y will act otherwise, in holding Y to this expectation, X will presume that Y will act on this reason and so ϕ . Thus this normative expectation entails a presumption about outcome that has the same rationalizing effect as the belief that Y will ϕ : it renders reliance willing.

Two-place trust is then related to this ‘moralized’ three-place trust. It cannot consist in a normative expectation because it is not instantiated in a particular case of reliance. Rather, it is a belief, but a belief whose content is something like a generalised version of the presumption made in this instance of trust, namely: were one to rely on the trusted, they would be moved by this fact. That is, to say that X trusts Y is to say that X believes that, for some ϕ , were he to rely on Y ϕ -ing, this would motivate Y to ϕ . This is not to propose the rejected reduction since this belief *qua* attitude is different to, and not reducible to, the normative expectation had in three-place trust. The commonality is that in both cases there is, as Uslaner (2002, 18) says, “the belief that [the trusted] share[s] your fundamental moral values”.

Thus, and in sum, acts of trusting, which is to say willing acts of reliance, instantiate four distinct kinds of trust, which for the purposes of this paper can be grouped to give three kinds. First, there is trust that is the conjunction of reliance with a belief about outcome. Given that we can ‘trust’ in this way without having any positive view of the ‘trusted’ party’s motivations – e.g. I might trust you to pay me because I know you fear the consequences of not doing so – I’ll call this *predictive trust* to make clear that trust is little more than a prediction about outcome (see (Faulkner 2007b, 880) and (Hollis 1998, 10)). Second, there is trust that is the conjunction of reliance with a normative expectation of outcome, which can be grouped with two-place trust since both are grounded by the belief that one’s relation to the trusted is normatively structured. As such,

I'll refer to both as *affective trust* (rather than 'moralistic' trust) to make it clear that trust is associated with the reactive attitudes — this is the kind of trust that can be betrayed — and let the context make clear whether it is two-place or three-place trust that is the concern (see (Faulkner (2007b, 881)). To these two attitudes of trust must be added a third: a generalized attitude of trust — trust with no specific object, one-place trust. Since it involves no particular belief about the trusted, this generalized form demonstrates the belief that our reliant interactions are normatively structured — that there is a right, or trustworthy, thing to do — combined with the optimistic view that others can generally be relied on to do this right, or trustworthy, thing. I'll refer to this as *generalized trust*. This attitude of generalized trust then supports a willingness to rely on others, where such specific acts of reliance will then be acts of affective trust. As such, affective and generalized attitudes of trust can be grouped by their shared *normative and optimistic* character. While predictive and affective trust can be grouped as *particularized attitudes* of trust: their object being some person or trust situation.³

³ Uslaner opposes particularized trust to generalized trust in terms of their psychological attitude, see (Uslaner 2002, 32-38). But this is because the *only* particularised attitude he considers is predictive trust. With this restriction in place, the contrast follows. However, once it is recognized that particularized attitudes of trust include affective trust, this contrast disappears.

The question now is the location of these kinds of trust within the polity. The primary instantiation of trust is in interpersonal contexts. A situation where one subject X anticipates relying on another Y, or does rely on Y in some particular way; these are situations where X might trust Y, or trust Y to ϕ . With respect to trust within the polity, there then seems to be two kinds of interpersonal engagement where trust might be found. First, trust can be horizontal: it can be between citizens; say, citizen X trusting citizen Y to ϕ , where ϕ -ing might be some transactional activity like the handing over of goods. Second, trust can be vertical: it can be between citizens and governors; say, citizen X trusting a government representative Y to ϕ , where ϕ -ing is an action prescribed by the representative's role in government. Then in addition to these interpersonal contexts there is the question of whether citizens approach each other and governors, or government, with a generalized attitude of trust. This identifies four possibilities of trust, so consideration of the social and political role of trust requires considering these four different domains, which I do in the next four sections.

2. Particularized Trust Between Citizens

Consider trust as it is found in daily interactions between citizens — horizontal trust. It is the multitude of these cooperative engagements that are the fabric of a society. Many of these interactions will be between parties that know one another — spousal sharing of childcare

arrangements for example. But it is a feature of complex modern societies that cooperating parties are increasingly unknown to one another — as is the case when one first hands one's child over at day-care. But in both cases cooperation only happens when there is a willingness to depend and this can be sustained by both predictive trust and affective trust. What the shift to a complex modern society would then seem to do is reduce the grounds that are available for predicting cooperative behaviour. After all, insofar as one knows one's interaction partner, one has a repository of reasons for predictive trust, so remove this knowledge and one's grounds for prediction could only seem to go down. Thus one might think, as Hobbes (1651) does, that it is one of the chief roles of the state to institute mechanisms, such as sanctions, that supply reasons for predictive trust. And one might think this not merely because an absence of grounds for predictive trust would dissolve the fabric of society, but also because it is natural to think that a greater level of predictive trust would be better for society since the more easily people can believe others can be relied on, the easier it will be to rely on them. In economic terms, improved predictive trust reduces transaction costs, which in turn are the biggest determiner of economic performance (see (Casson 1991, 1)). However, things are not quite so straightforward because mechanisms designed to promote predictive trust can in fact undermine affective trust. Thus, O'Neill (2002, 73) argues that the institutions put in place to ensure transparency, to combat the perceived decline in trust, are often counter-productive and undermine trust. Braithwaite and Makkai (1994)

illustrate this with a detailed case study of how the introduction of quality care standards into an Australian nursing home undermined the achievement of the standards it was designed to monitor. The distinction between the two kinds of particularized trust — predictive and affective — then offers some theoretical traction on why this might be the case.

Predictive trust is supported by whatever empirical ground one has for predicting that another will prove reliable. Broadly, there are two distinct kinds of grounds, which might be labelled *game theoretical* and *trust-based*, where both focus on the motivations that someone has to prove reliable. Game-theoretical grounds focus on the instrumental reasons a person can have to prove reliable. Here Pettit (2002) observes that an act of reliance can itself supply such a reason, which he calls a *trust-responsive reason*, since we are sensitive to the positive status implicit in the evaluation that we are reliable, an act of reliance can give us a reason to be reliable. These game-theoretical grounds form part of the general empirical theory each of us has about the motivations of others and the operation of the world generally. As such, there is no fundamental shift in perspective between predictions which are so based and the general predictions we make about the physical world, such as that a branch might hold our weight, for instance. In both cases one takes an objective, or ‘scientific’, approach to matters.

However, prediction can also be grounded on trust-based reasoning. Identifying this class of reasons requires more than a consideration of motivation, it requires *empathy*, or a capacity to imagine motivations as subjectively held. Thus, I can be confident that you will turn up because I imagine you imagining my waiting and in imagining this presume you will be moved by it. However, where prediction has this ground, it would be wrong to describe trust as predictive: in reasoning in this way, one has entered a normative domain where blame and resentment can be appropriate responses — for instance if you leave me waiting with no apology; so the trust in play is affective. The slightly complex relationship between these two kinds of trust then lies in the interaction of the different grounds these attitudes are held on, and it may be traced in the following points.

(1) *Affective trust is sufficient for predictive trust, other things being equal.* If I presume you are motivated by my relying on you to do something, then, other things being equal, I will predict you will do this thing. And if other things are not equal, so that I cannot make this prediction, then the chances are that I won't be able to affectively trust you in this regard, and as such won't make this presumption either. (However, see point (3) below.)

(2) *Affective trust is necessary for cooperation generally.* If the only ground I have for thinking that you will prove reliable is a judgement of particular motivating facts, then I have no

reasons, or no good reasons, for thinking you will prove reliable if these facts cease to hold. But that they can do so is just the ‘free-rider’ worry: why should someone do something that is of some cost to them when they can get away with not doing so? Thus the only way to ensure cooperation in all circumstances would be to find the motivation to be cooperative in the facts that must hold in every case. And this is just what is expected in affective trust: that the trusted be motivated by the fact of the trusting party’s reliance. Thus cooperation can only be generally secured to the extent that there is affective trust and the correlative kind of trustworthiness (I elaborate this argument in the next section, §3).

(3) *Predictive trust is not necessary for affective trust.* It is possible to give people the benefit of the doubt. Even if the evidence suggests someone is motivated to defect, it is possible, on many if not all occasions, to bracket this evidence and choose to operate on the presumption that they will nevertheless see one’s relying on them in the ‘right’ light (It is this possibility that allows rehabilitation, which is what Holton (1994, 63) imagines in giving a case where a shopkeeper employs an ex-convict and affectively trusts him with the till.)

(4) *There is a tension between reasoning game theoretically and engaging in trust based reasoning.* This follows from (3). Suppose the ex-convict offered the reassurance that the shopkeeper has “nothing to worry about”. In saying this, there is an implicit demand for affective

trust, which is the demand that the shopkeeper believes this *despite* the evidence. But the tension here is general. Thus, compare the case where the shopkeeper similarly trusts a close friend. And suppose the friend offers the same reassurance. In this case, the implicit demand would be that the shopkeeper believes the reassurance *because of affective trust rather than* because of the evidence. And in both cases, were the shopkeeper to turn on the CCTV to “keep an eye on things”, this act would be taken as non-trusting even though it might be support prediction. Thus we prefer our engagements to be trust-based since it credits our agency, which is denied when our actions are viewed from an objective or ‘scientific’ perspective.

(5) *Predictive trust and affective trust are mutually reinforcing.* The more confident one is that someone can be relied on, the easier it is to affectively trust them; and, conversely, the less confident, the harder it is to affectively trust, where the limiting case is that if one knows that they will prove unreliable, one will not be able to presume they will be moved by one’s dependence, so one will not be able to expect this of them.

It is (4) – the tension between predictive and affective trust – I suggest, that underlies the empirical data — observed by (Braithwaite and Makkai 1994) and theorised by (O’Neill 2002) — showing that if you treat people as knaves, they are more likely to be so. What explains this data is that we want others to treat as agents and not objects. However, from (5) — that predictive and

affective trust operate synergistically — there is reason for government to implement mechanisms whose operation give grounds for predictive trust, such as sanctions for defection, since doing so should increase levels of trust. Moreover, from (1) — that affective trust can suffice for predictive trust — it follows that trust can beget trust: there can be a virtuous circle with an increase in predictive trust enabling more affective trust which then supports further predictive trust. And from (2) — that trust is necessary for cooperation generally — it seems that there is no legislative option: something must be done to create an environment where trust can develop. But then how is this to be done when the legislative mechanisms that support predictive trust do so by reducing the gains of knavery, and thereby treat people as knaves?

This is a difficult question. O'Neill (2002, 133) proposes that rather than focus on transparency, and thereby supplying grounds for prediction, government should aim to improve accountability, and openness in public life. Meanwhile Braithwaite (1998, 351) proposes *dynamic regulatory institutions*, which follow “the regulatory pyramid”.

Trust first, and thereby get the efficiency benefits of trust in most cases, but motivate trust as obligation [affective trust] by signalling very clearly a preparedness to escalate intervention to progressively less trusting interventions when trust is abused. (Braithwaite 1998, 352).

This proposal, I think, exactly fits our thinking about trust in its two particularized forms.⁴

3. Generalized Trust Between Citizens

Consider now the other horizontal dimension of trust: whether citizens' attitude towards each other is one of generalized trust? The shift in focus here is from a given citizen X trusting another citizen Y to ϕ — or trusting Y generally — to whether X is generally trusting towards other citizens; that is, is X's background attitude one of trust? Standardly, generalized trust has been measured by way of the survey question, 'Generally speaking would you say that most people can be trusted or that you can't be too careful in dealing with people?' This question has formed part of three big longitudinal social studies that have thereby mapped levels of generalized trust.⁵ The levels thus indicated have been found to correlate with civic participation (Brehm and Rahn 1997) and well-functioning institutions (Delhey and Newton 2005), and to inversely correlate with corruption (Uslaner 2013), violent crime (Lederman, Loayza, and Menéndez 2002) and income inequality (Knack and Zak 2002). Moreover the cited claims that trust furthers economic prosperity and growth (Putnam 1993 and Fukuyama 1995) and fosters good government (Knack 2002) are claims

⁴ This proposal is also similar to tit-for-tat: cooperate in the first round but retaliate to any non-cooperation, which also does well at securing cooperation. See (Axelrod 1984).

⁵ The surveys being: American General Social Survey, World Values Survey, and European Social Surveys.

made about generalized trust. These studies show the social and political importance of generalized trust; and this importance could also be reached by way of a philosophical argument (which elaborates the argument given for claim 2 in the last section).

Suppose that our trust in others is no more than predictive. And consider the schematic case of X trusting Y to ϕ . Cooperation is then possible when Y possesses reasons to ϕ , and X possesses grounds for believing this. Hardin's encapsulation of interest is one such reason and ground. But there are more, and Williams (1988, 118) identifies four general reasons for cooperation, which are also grounds for predicting it; these are: fear of sanctions, particular self-interest, a positive evaluation of cooperation, and a positive evaluation of friendly relations. All these reasons for cooperating are instrumental, including the latter positive evaluations, because, on this conception of trust, value is conceived as utility — as a good to be weighed and calculated, like money. Thus, a positive evaluation of friendly relations, say Y's desire to maintain good relations with X, can be weighed against particular self-interest, say Y's desire not to ϕ . This is equally true of Pettit's (2002) trust-responsive reason: the desire for the status associated with proving reliable will be offset against the costs of proving so. However, insofar as these reasons are instrumental, they are not all on a par: there is one — particular self-interest — that is operationally guiding. And the problem is this reason does not favour cooperation; it is cooperation neutral: the issue in every case will be the costs and benefits for Y in ϕ -ing. Insofar as we then

credit others with this constant pull, there will be the tendency for cooperation to breakdown, or not get started. For the worry is that others will listen to the counsel of Hobbes's Foole who whispers the irrationality of honouring an agreement made with another who has already fulfilled their part. Hume's farmers, each of whom needs the other's help to harvest their corn, illustrate this worry. Knowing that the other, when their corn is in the barn, will feel the pull of labouring no further, neither is willing to be the first to help the other. The consequence is that "both ... lose our harvests for want of mutual confidence" (Hume 1740, 521).

Williams's list of reasons for predictive trust names the various ways this worry can be addressed and cooperation grounded. For instance maybe the farmers care to maintain friendly relations because they have an eye on future harvests (not merely next year's harvest or the next few years but the indeterminate future, otherwise a backwards induction would start). But the point is not that the farmers cannot cooperate, but that all the while self-interest is the only constant, cooperation can be fragile. Thus, following Hobbes (1651), one might think that it is one of the chief roles of the state to institute mechanisms, such as sanctions, that supply reasons for predictive trust. However, while such mechanisms will help secure cooperation, the structural problem remains: no mechanism will be perfect, so the Foole's advice is always there (see (Hollis 1998)). If one can get away with defecting, why not do so? The general issue is that all the while our

grounds for predictive trust are limited to particular reasons for predicting things will work out, we will frequently have insufficient reason to cooperate given the known pull of self-interest.

Nevertheless, cooperation is not a rarity and is ordinarily, if not always, secured. The reason for this, I suggest, is that we take the trust situation to be normatively structured. This argument (which repeats that given in (Faulkner 2017)) supports Uslaner's (2010, 112) claim that what is needed for cooperation more generally is *moralistic* trust, or the idea that there is "a moral dimension" to our interactions. In presuming that our interactions are normatively structured, or have a moral dimension, we take the trusted to have some reason to cooperate *in every case*. This reason, though defeasible, gives a constant pull towards cooperation rather than orthogonal to it. There is a shift in thinking here: away from predictive trust and towards affective trust, whose expectation is not merely subjective – a belief about outcome – but normative. The content of this normative expectation, I've suggested, is that the trusted party will take one's reliance as a reason for doing what one relies on them doing. To say that citizen X is trusting, or has a generalized attitude of trust, is then to say that X is willing to affectively trust. To explain this willingness, the attitude of generalized trust must then be understood as the conjunction of a belief that the reliant interactions are often normatively structured with the optimistic view that people can be relied on to do the 'right thing'. It is this conjunction that supports the presumptions made in a given instance

of affective trust, which rationalize cooperation. And it is this generalized attitude of trust that must be widespread for cooperation to be equally so.

In the political domain, Lenard (2008) argues that in order for democratic institutions to function, there must be widespread voluntary compliance with their rules and regulations. This requires that citizens have the general expectation that others will comply, which, for the reasons just adumbrated, requires generalized trust. Thus, generalized trust “helps build large-scale, complex, interdependent social networks and institutions” (Warren 1999, 9). Three further consequences follow. First, given that generalized trust is simply a willingness to affectively trust, a consequence of this is that one cannot persuade someone to be trusting, but would expect their disposition to be so, to be a feature of their make-up or character. Thus, this disposition, Uslaner (2002, 26) argues, is something that is developed “early in life”. Since these individual dispositions would then be stable over time, one would expect societal levels of generalized trust, which measure the aggregation of these dispositions, to be similarly stable. And this is indeed what has been found (see (Bjørnskov 2007)). Second, a willingness to affectively trust presumes that the trusted will see the interaction situation in the same light — as structured by the same normative reasons. It follows that “the standard survey question is thus a measure of how widely people view their moral community” (Uslaner 2002, 28). Third, while it is hard to positively change a person’s disposition to trust and be trustworthy, it is possible and easier to negatively change it: an abuse

of trust, or even, as observed in the last section, no more than the manifestation of the view that someone is a knave can be sufficient to make this so. Thus, while there is much stability in levels of generalized trust, surveys of trust have also notably found levels of trust to have been falling in the States since the 1970s.⁶

4. Particularized Trust Between Citizens and Government

Consider trust as it is found between citizens and governors, or government representatives — vertical trust. That is X trusting government representative Y to ϕ , or trusting Y generally. Given his understanding of particularized trust as encapsulated interest, Hardin (2002) concludes that it is not possible for citizens to trust government representatives. Trust only exists in relationships that allow a judgement of encapsulated interest, and this is not one of those. Hence Hardin (2002, 152) concludes that “trust in government is conceptually and epistemologically impossible for most citizens in large modern societies”. Pettit (1998) makes a similar claim: with respect to its vertical axis, the grounds that a citizen has for particularized trust are very limited, and these are mainly trust-responsive grounds. Both these claims, in my view, are incorrect: citizens do have

⁶ See (Putnam 2000) The explanation of this is no doubt multi-factorial, but one would expect it to refer to social changes that disrupt the sense of moral community — such as the increasing wealth inequality. See (Uslaner 2010)

available grounds for particularized trust, and grounds of trust-responsiveness are no more available than those of encapsulated interest.

Pettit's (1998) argument start from a distinction between 'personal' and 'impersonal' trust: a citizen's trust in government representatives cannot be 'impersonal' because we do not have the grounds needed for such trust. What would be needed for impersonal trust is for government representatives to be constrained in their operative roles and for citizens to know this to be so. However, Pettit argues, it is implausible to suppose that there are any constraints rigid enough to support this judgement.

No matter how restricted the tasks assigned to government, there are bound to be areas of significant discretion that are left to legislative, executive and judicial authorities. ... [T]here is no possibility of constraining government agents to more or less uniquely determined choices. (Pettit 1998, 301-2)

What follows from this is that citizens are forced "either to trust the authorities on a personal bias [*sic.*] or to distrust them on a personal basis" (Pettit 1998, 300). The grounds for 'personal' trust are then two-fold. First, they are "provided by the fact that the authorities are virtuous and trustworthy" (Pettit 1998, 306). Second, they are provided by the thought that the authorities are *trust-responsive*: that they want to be seen to be virtuous and trustworthy, and will, as such, respond in such a manner to acts of trust. Now I argued above (in §2) that this trust-responsive

mechanism gives ordinary empirical ground for predictive trust — it being an instance of game theoretical reasoning, or reasoning as to interest. However, the thought that the authorities are virtuous rather suggests a trust-based ground, which would imply the quite different affective conception of trust. Moreover, Pettit seems to affirm this implication when clarifying the nature of trust-responsive reasons.

These extra reasons that people have for trusting those in power are not reasons of trustworthiness ... but reasons of trust responsiveness. They come of the belief that even if the agents are not moved by the fact of other's relying on them, in the manner of truly virtuous and trustworthy individuals, they will at least be moved by the fact that those others will think well of them for proving reliable and will think badly of them for proving unreliable. (Pettit 1998, 308)

The crucial matter is then the place of the thought that the trusted — government representative — will be moved by the fact of others relying on them. If this is an expectation that the trusting party — citizen — has *of the trusted*, if it is something the trusted is held to in trust, then 'personal trust' is what I've been call affective trust. However, if this is merely a belief *about the trusted*, then it does no more than name one of Williams's (1988) sources of motivation: a positive valuation of cooperation. Pettit's intention, I think, is the latter — and if it were the former it is not clear how personal trust could have both the grounds he suggests. As such, this distinction between 'impersonal trust' and 'personal trust' is a distinction within what I have been calling *predictive*

trust; it amounts to a distinction within the grounds one can have for prediction: either these grounds involve some reference to one's reliance — in which case the trust is *personal* — or they do not, in which case it is *impersonal*.

When the concern is trusting a government representative, this focus on predictive trust is the right one to have. For consider what is demanded in affective trust – say by reference to the expectation I have when I wait for you in the café. The expectation I hold you to is that you will do something (arrive in a timely fashion) just because I rely on your doing this (and take you to know this, given our arrangement). Schematically, I expect you to ϕ because I rely on you ϕ -ing. But to suppose that a government representative should respond, in this way to one's particular dependence, is to suppose them *corrupt*. The response of any government representative needs to be third personal not second personal; that is, it should not be a response to the *fact* of the citizen's reliance but to *facts about* the citizen's reliance; that is, to the details of the citizen's case, where these are considered in abstraction from the fact that the citizen relies in the way that they do. To suppose otherwise is to suppose that citizens have the capacity to give government representatives reasons to act in a way that undermines whatever government institution they represent.

So Pettit is correct to focus solely on predictive trust. However, the worry just raised about affective trust might equally be raised about personal trust — or predictive trust where one's

grounds for prediction involves some reference to the fact of one trusting. Thus, consider a case where personal trust is grounded by trust responsiveness. If the supposition is really that a government representative will be responsive to my particular citizen trust — if it is that they want the kudos that goes with being seen to do the ‘right’ thing — then the supposition is that I have the individual capacity to influence their actions, and this is just to suppose that they fail in their governing role. These grounds should be no more available than a calculation of what lies in the government representative’s interests: their personal interests should not figure in our relations.

Now consider the ground for predictive trust supplied by something impersonal like a belief in the trusted’s virtue. If it is assumed, as I think it should be, that this is *not* the ground that one has in affective trust, and one is not expecting a government representative to respond to the fact *of one’s reliance*, but to facts *about one’s reliance*, then it seems that the trust is straightforwardly predictive trust of Pettit’s impersonal type. But this kind of trust in government Pettit takes to be impossible? Given that the belief that a government representative is virtuous would seem to ground it, it is worth revisiting Pettit’s reasons for this impossibility claim.

The reason one lacks impersonal grounds for trusting government representatives, according to Pettit (1998, 301) is that these agents “face relatively unconstrained choices”. The point here seems to be no more than that government representatives *always have some degree of*

discretion in deciding how to act. But this is true in every case of trust. No matter how good one's ground for predictive trust, it will never be such as to guarantee that the trusted acts in the way one relies on them to act. Were one's ground, *per impossible*, to guarantee this, it would be wrong to say one *trusted*. But, in fact, the most that one can have is inductive grounds: some kind of generalization of, or judgement from, experience. And here Hardin (2002, 159) agrees: we can have *confidence* in government, which is based on inductive grounds. But then it follows that there is no reason to say we do not have impersonal grounds to be confident of government representatives. And that they can be virtuous, or be "sufficiently conscientious about their brief", is a fact that, if we believe it to hold, gives such an impersonal ground. It is then such grounds for predictive trust that are wanted, and not grounds for trust (be it affective or personal). Moreover, it is these kinds of grounds for predictive trust in itself that government then seeks to supply with the establishment of institutions of transparency (again see (O'Neill 2002, 73)).

5. Generalized Trust Between Citizen and Government

In the last section (§4), I argued that it would be wrong for citizens to affectively trust government representatives. This argument considered an individual citizen's relation to a particular government representative. But what of a citizen's relation to government representatives in

general, which might be put as the citizen's relation to *the* government? Here there is an interesting split between affective trust and generalized trust.

With respect to affective trust, an individual citizen cannot affectively trust the government, not because it would be wrong to do so, but because the government is not an appropriate object of affective trust. One can predict — so trust in this sense — that one's car will start. But one cannot affectively trust one's car to start. One cannot do so because one's car cannot be correspondingly trustworthy: it cannot start for the reason that one needs it to. So one cannot be betrayed by one's car's failure to start.⁷ Similarly, one cannot affectively trust the government because it likewise does not have the agency necessary for assigning due deliberative weight to the fact of the trusting party's dependence. It cannot reflectively engage in trust-based reasoning. This is not to say that the government cannot be regarded as an agent rather it is to say that it does not have the kind of self-conscious agency that is needed to be the appropriate object of affective trust. (In this respect, I would echo Hardin's (2002, 152) claim and claim that *affective* "trust in government is conceptually ... impossible".) And insofar as affective trust is ruled out, one would expect generalised trust to be problematic, given that it is a willingness to affectively trust. However, what rationalizes a given instance of affective trust is the belief, or presumption, that

⁷ Hence the comedy in: <<https://www.youtube.com/watch?v=mv0onXhyLIE>> (accessed 4/11/16).

one shares the same ‘moral universe’ with the trusted, combined with the optimism that the trusted will do the ‘right thing’. And while generalized trust can underwrite a willingness to affectively trust, it does so because it is to be identified with this attitudinal background. As such it is perfectly appropriate to talk of generalized trust of government. What is expressed here is just the idea that we can normatively expect government, and its representatives, to do the right thing — to make good decisions on our behalf, decisions that take account of our needs and dependencies — and are optimistic in holding this normative expectation. (It is unproblematic if this attributes a degree of agency to government: we happily talk of the government paying, or cutting, welfare, offering tax breaks, convicting the offender and so on, see (List and Pettit 2011)).

When a generalized trust of government is conjoined with generalized trust between citizens, there is, what Lenard (2015, 353) calls, “a climate of trust”. A climate of trust produces a state with many virtues, principally because it ensures that citizens willingly comply with government legislation. These virtues, Lenard (2015, 353) states, “range from the simple willingness to abide by laws that cannot, always, be coercively enforced; to high rates of participation in democratic decision-making procedures; to a widespread commitment to contribute to social-justice policies.” And more abstractly, widespread “generalized trust is connected to a number of dispositions that underwrite democratic culture, including tolerance for pluralism and criticism” (Warren 1999, 9).

Of course, insofar as generalized trust is identified, in part, as an optimistic attitude that government will do the right thing, it ought to have the same correctness conditions as belief. So that one can have a generalised attitude of trust towards government, and that this climate of trust fosters various state virtues, does not imply that it is right to have this attitude. Rather, generalized trust will be right to the extent that there are grounds for it, which is to say evidence that the government will act in ways that satisfy our normative expectations of it. When the government does so, it makes sense to talk about responsible or ‘trustworthy’ government. The existence or not of evidence of responsible government is then an empirical matter.

Here two empirical observations might be expected. First, we are likely to have expectations of government when it comes to our welfare — just as we are more likely to rely on a friend for a loan to pay the rent than for entertainment. Interactions where our welfare is at issue are more likely to be normatively structured. It follows that levels of generalized trust are likely to be higher when governments institute welfare programmes. Thus Rothstein and Uslaner (2005) argue that universal redistribution policies foster generalized trust. Second, observation of corruption undercuts one’s reason for thinking that the government will do the right thing. It follows that levels of generalized trust are likely to be inversely correlated with believed levels of corruption. This is also argued by Rothstein and Uslaner (2005) (and see (Uslaner 2013)). Relatedly, Maclachlan (2015) shows that when there has been some breach of the expectations

associated with generalized trust, public apology can contribute to the sense of community between persons, and so preserve generalized trust. And when generalized trust is inappropriate, what should it be replaced by? Some suggest distrust — (Allard-Tremblay 2015) and (Krishnamurthy 2015) — Lenard (2008) suggests mistrust.

6. Conclusion

In order to understand the role of trust in society it is necessary to proceed from a theory of trust, or so this paper has proposed. It thereby starts from a pre-existing theory of trust, which advances two distinctions: attitudes of trust can take different objects and be of different psychological kinds. These two distinctions generate three distinct attitudes of trust: predictive trust, affective trust, and generalized trust. The first two of these attitudes can be grouped together as particularized attitudes of trust in that they are focused on some person or trust situation. And the latter two attitudes of trust can be grouped together as normative forms of trust. With respect to the political domain, there are then two locations where trust is relevant and can play a social and political role. There is the trust that is found between citizens, where this locates trust along a horizontal political axis. And there is the trust that is found between citizens and government representatives, where this locates trust along a vertical political axis. Applying the two categories — particularized and

generalized trust — to these two axes then generates four domains where some account of the role of trust is needed. In brief and to recap, this paper has suggested the following. With respect to particularized attitudes of trust between citizens, there is a tension between predictive and affective trust that is politically important: the pursuit of grounds for predictive trust can undermine the possibility of affective trust. With respect to the generalized attitude of trust that citizens can have towards each other, this attitude is fundamental in that no trust, and so no society, is possible unless this attitude is secured first. With respect to particularized attitude of trust between citizens and government representatives, the only trust that is appropriate is predictive, and this can only have certain grounds. And with respect to the generalized attitude of trust that a citizen can have towards government, this secures a range of virtues for the state and so should be cultivated and preserved where possible. At least, this is a sketch of the conclusions reached; obviously, much more was said of each, and there is much more that can be said. A further principle conclusion of the paper is simply the trust schema advanced — the identification of these four domains of trust and the categorisation of trusting attitudes — where this schema itself has consequences. Thus, and as stated in the introduction, it implies that Hardin is wrong to say that trust in government is impossible: one can both predictively trust government representatives and have a generalized attitude of trust towards government. And Uslaner is wrong to separate the attitude of trust from

the attitudes that citizens can have towards government: the 'moral' dimension of trusting other persons can equally apply to the attitude of trust we can have towards government.

Allard-Tremblay, Y. 2015. "Trust and Distrust in the Achievement of Popular Control". *Monist* 98 (4):375-390.

Axelrod, R. 1984. *The Evolution of Cooperation*. New York: Basic Books.

Baier, A. 1986. "Trust and Antitrust". *Ethics* 96:231-60.

Becker, Lawrence. 1996. "Trust as Noncognitive Security About Motives". *Ethics* 107:43-61.

Bjørnskov, C. 2007. "Determinants of Generalized Trust: A Cross-Country Comparison". *Public Choice* 130:1-21.

Braithwaite, John. 1998. "Institutionalizing Distrust, Enculturating Trust". In *Trust and Governance*, edited by M. Levi and V. Braithwaite. New York: Russell Sage Foundation.

Braithwaite, John, and T. Makkai. 1994. "Trust and Compliance". *Policing and Society* 4:1-12.

Brehm, J., and W.M. Rahn. 1997. "Individual-Level Evidence for the Causes and Consequences of Social Capital". *American Journal of Political Science* 41:999-1023.

Casson, Mark. 1991. *The Economics of Business Culture*. Oxford: Clarendon Press.

Darwall, Stephen. 2017. "Trust as a Second-Personal Attitude (of the Heart)". In *The Philosophy of Trust*, edited by P. Faulkner and T. Simpson. Oxford: Oxford University Press.

Delhey, J., and K. Newton. 2005. "Predicting Cross-National Levels of Social Trust: Global Pattern or Nordic Exceptionalism?". *European Sociological Review* 21:311-27.

Faulkner, Paul. 2007a. "A Genealogy of Trust". *Episteme* 4 (3):305-321.

———. 2007b. "On Telling and Trusting". *Mind* 116 (464):875-902.

- . 2010. "Norms of Trust". In *Social Epistemology*, edited by D. Pritchard, A. Haddock and A. Millar. Oxford: Oxford University Press.
- . 2011. *Knowledge on Trust*. Oxford: Oxford University Press.
- . 2012. "The Practical Rationality of Trust". *Synthese*:1-15.
- . 2014. "The Moral Obligations of Trust". *Philosophical Explorations*.
- . 2015. "The Attitude of Trust Is Basic". *Analysis* 75 (3):424-429.
- . 2017. "The Problem of Trust". In *The Philosophy of Trust*, edited by P. Faulkner and T. Simpson. Oxford: Oxford University Press.
- Fukuyama, F. 1995. *Trust: The Social Virtues and the Creation of Prosperity*. New York: Free Press.
- Gambetta, D., ed. 1988. *Trust: Making and Breaking Cooperative Relations*. Oxford: Basil Blackwell.
- Glaeser, E., D. Laibson, J. Scheinkman, and C. Soutter. 2000. "Measuring Trust". *Quarterly Journal of Economics* 113:811-845.
- Govier, Trudy. 1992. "Trust, Distrust, and Feminist Theory". *Hypatia* 7 (1):16-33.
- Hardin, R. 2002. *Trust and Trustworthiness*. New York: Russell Sage Foundation.
- Hawley, Katherine. 2014. "Trust, Distrust and Commitment". *Nous* 48 (1):1-20.
- Hobbes, T. 1651. *Leviathan*. 1968 ed. Harmondsworth: Penguin.
- Hollis, M. 1998. *Trust within Reason*. Cambridge: Cambridge University Press.
- Holton, R. 1994. "Deciding to Trust, Coming to Believe". *Australasian Journal of Philosophy* 72 (1):63-76.
- Hume, D. 1740. *A Treatise of Human Nature*. Edited by Selby-Bigge. Oxford: Clarendon Press.
- Jones, K. 1996. "Trust as an Affective Attitude". *Ethics* 107 (1):4-25.
- Knack, S. 2002. "Social Capital and the Quality of Government: Evidence from the United States". *American Journal of Political Science* 46:772-85.
- Knack, S., and P.J. Zak. 2002. "Building Trust: Public Policy, Interpersonal Trust, and Economic Development". *Supreme Court Economic Review* 10:91-107.

- Kreps, D.M., P. Milgrom, J. Roberts, and R. Wilson. 1982. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma". *Journal of Economic Theory* 27 (2):245-252.
- Krishnamurthy, M. 2015. "(White) Tyranny and the Democratic Value of Distrust". *The Monist* 98 (4):391-406.
- Lederman, D., N. Loayza, and A.M. Menéndez. 2002. "Violent Crime: Does Social Capital Matter?". *Economic Development and Cultural Change* 50:509-39.
- Lenard, Patti. 2005. "The Decline of Trust, the Decline of Democracy?". *Critical Review of International Social and Political Philosophy* 8 (3):363-378.
- . 2008. "Trust Your Compatriots, but Count Your Change: The Roles of Trust, Mistrust and Distrust in Democracy". *Political Studies* 56:312-332.
- . 2015. "The Political Philosophy of Trust and Distrust in Democracies and Beyond". *The Monist* 98 (4):353-359.
- List, C., and P. Pettit. 2011. *Group Agency*. Oxford: OUP.
- Maclachlan, Alice. 2015. "'Trust Me, I'm Sorry': The Paradox of Public Apology". *The Monist* 98 (4):441-456.
- McGeer, Victoria, and Philip Pettit. 2017. "The Empowering Theory of Trust". In *The Philosophy of Trust*, edited by P. Faulkner and T. Simpson. Oxford: Oxford University Press.
- Nannestad, Peter. 2008. "What Have We Learned About Generalized Trust, If Anything?". *Annual Review of Political Science* 11:413-436.
- O'Neill, O. 2002. *A Question of Trust*. Cambridge: Cambridge University Press.
- Pettit, Philip. 1998. "Republican Theory and Political Trust". In *Trust and Governance*, edited by V. Braithwaite and M. Levi. New York: Russell Sage Foundation.
- . 2002. "The Cunning of Trust". In *Rules, Reasons, and Norms*, edited by P. Pettit. Oxford: Clarendon Press.
- Putnam, Robert. 1993. *Making Democracy Work. Civic Traditions in Modern Italy*. Princeton, NJ: Princeton University Press.
- . 2000. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster.

- Rothstein, B., and Eric Uslaner. 2005. "All for All - Equality, Corruption, and Social Trust". *World Politics* 58:41-72.
- Uslaner, Eric. 2010. "Trust and the Economic Crisis of 2008". *Corporate Reputation Review* 13 (2):110-123.
- . 2013. "Trust and Corruption Revisited: How and Why Trust and Corruption Shape Each Other". *Quality and Quantity* 47 (6):3603-3608.
- Uslaner, Eric M. 2002. *The Moral Foundations of Trust*. Cambridge: Cambridge University Press.
- Warren, M. 1999. "Introduction". In *Democracy and Trust*, edited by M. Warren. Cambridge: Cambridge University Press.
- Williams, B. 1988. "Formal Structures and Social Reality". In *Making Sense of Humanity and Other Philosophical Papers 1982-1993*, edited by B. Williams. Cambridge: C.U.P.
- . 2002. *Truth and Truthfulness*. Princeton: Princeton University Press.