

Philip Pettit

Trust, Reliance and the Internet

Abstract: Trusting someone in an intuitive, rich sense of the term involves not just relying on that person, but manifesting reliance on them in the expectation that this manifestation of reliance will increase their reason and motive to prove reliable. Can trust between people be formed on the basis of Internet contact alone? Forming the required expectation in regard to another person, and so trusting them on some matter, may be due to believing that they are trustworthy; to believing that they seek esteem and will be rationally responsive to the good opinion communicated or promised by an act of trust; or to both factors at once. Neither mechanism can rationally command confidence, however, in the case where people are related only via the Internet. On the Internet everyone wears the ring of Gyges; everyone is invisible, in their personal identity, to others.

1. Introduction

Words like “trust” and “reliance” are used as context requires, now in this way, now in that, and they serve to cover loose, overlapping clusters of attitudes and actions. Here I invoke some theoretical license, however, and use the terms to tag distinct phenomena: ‘reliance’ a generic phenomenon, ‘trust’ a species of that genus. I want to argue that while the Internet may offer novel, rational opportunities for other forms of reliance, it does not generally create such openings for what is here called trust.

The paper is in three sections. In the first, I set up the distinction between trust and reliance. In the second, I outline some different forms that trust may take. And then in the final section I present some reasons for thinking that trust as distinct from other forms of reliance is not well served by interactions on the Internet, at least not if the interactants are otherwise unknown to one another. (The paper is a follow-up to Pettit 1995, and draws freely on some arguments in that piece.)

The Internet is exciting in great part because of the way it equips each of us to assume different personas, unburdened by pre-given marks of identity like gender, age, profession, class, and so on. A very good question, then, is whether people can develop trust in one another’s personas under the shared assumption that persona may not correspond to person in such marks of identity. Suppose that you and I appear on the Internet under a number of different names, developing a style that does with each: I am both Betsy and Bob, for example, you are both Jane and Jim. The question is whether you as Jane can trust me as Bob, I as Betsy can trust you as Jim, and so on. But good though it is, I should stress

that this is not the question that I try to deal with here (Brennan/Pettit 2004b; McGeer 2004a). My focus is rather on how far real-world, identity-laden persons may achieve trust in one another on the basis of pure Internet contact, not just how the Internet personas they construct can achieve trust in one another on that basis.¹

2. Trust and Reliance

Trust and reliance may be taken as attitudes or as actions but it will be useful here to take the words as primarily designating actions: the actions whereby I invest trust in others or place my reliance in them. So what then distinguishes relying on others in this sense from trusting in them?

To rely on others is just to act in a way that is premised on their being of a certain character or on their being likely to act under various circumstances in a certain way. It won’t do, of course, if the guiding belief about others is just that they have a low probability of displaying the required character or disposition. The belief that others will prove amenable to one’s own plans must be held with a degree of confidence that clearly exceeds 0.5. To rely on others, as we say, is to manifest confidence in dealing with them that they are of the relevant type or are disposed to behave in the relevant way.

I may rely on others in this sense in a variety of contexts. I rely on automobile drivers to respect the rules of the road when I step out on to a pedestrian crossing. I rely on my doctor’s being a competent diagnostician when I present myself for a medical examination. I rely on the police to do their duty when I report a crime I have just witnessed. In all of these cases, reliance is a routine and presumptively rational activity. If we are Bayesians about rationality, then we will say that such acts of reliance serve to promote my ends according to my beliefs and, in particular, that they serve to maximize my expected utility.

Relying on others in the sense exemplified here is not sharply distinguishable from relying on impersonal things: relying on the strength of the bridge, relying on the accuracy of the clock, and so on. True, the reliance is something I may expect those on whom I rely to notice but this does not appear to be essential. The important point in the cases surveyed is that relying on someone to display a trait or behaviour is just acting in a way that is shaped by the more or less confident belief that they will display it. And relying on a person in that sense is not markedly different from relying on a non-personal entity like a bridge or a clock or perhaps just the weather.

Acts of rational reliance on other people, such as our examples illustrate, do not count intuitively as acts of trust; certainly, they do not answer to the use of the word ‘trust’ that I treat here as canonical. Trusting someone in the sense

¹ While this question is more mundane than that which I ignore, it is in one respect of greater importance. A form of trust is intuitively more significant, the greater the potential gains and losses with which it is associated. And by this criterion person-to-person trust is likely to be of more significance than trust between personas. It may put one’s overall fortune at stake, where persona-to-persona trust will tend to involve only stakes of a psychological kind.

I have in mind—and it is a sense of trust that comes quite naturally to us—means treating him or her as trustworthy. And treating someone as trustworthy involves assuming a relationship with the person of a kind that need not be involved in just treating someone as reliable. To treat someone as reliable—say, as a careful driver, a competent doctor, a dutiful police officer—means acting on the confident belief that they will display a certain trait or behaviour. It would be quite out of place to say that whenever I treat a person as reliable in this way I treat them as trustworthy. Thus I might be rightly described as presumptuous if I described my attitude towards the driver or doctor or police officer as one of treating the person as trustworthy. The washerwomen of Koenigsberg might as well have claimed that they treated Kant as trustworthy when they relied on him for taking his afternoon walk at the same time each day.

The cases of reliance given, which clearly do not amount to treating someone as trustworthy, are all instances of rational reliance, as we noticed. Does this mean that when I go beyond mere reliance and actually trust a person—put my trust in the person—I can no longer be operating in the manner of a rational agent? Does it mean, as some have suggested, that trust essentially involves a leap beyond rationality, a hopeful but rationally unwarranted sort of reliance; if you like, a Hail-Mary version of the practice? This suggestion would leave us with a paradox that we might phrase as follows. If trust is rational then it is not deserving of the name of 'trust'—not at least in my regimented sense—and if it deserves the name of 'trust' then it cannot be rational.

Happily, however, there is an alternative to this suggestion, and a way beyond this paradox. The assumption behind the suggestion is that the only factor available to mark off ordinary reliance from trust is just the rationality of the reliance. But this is mistaken. The acts of reliance considered are distinguished, not just by being rational, but also by being, as I shall put it, interactively static. And what distinguishes trust from reliance is the interactively dynamic character of the reliance displayed, and not any necessary failure of rationality. So at any rate I argue.²

My relying on others will count as interactively dynamic when two special conditions are fulfilled; otherwise it will be static in character. The first condition required is that the people on whom I rely must be aware of that fact that I am relying on them to display a certain trait or behaviour; that awareness must not be left to chance—in the paradigm case, indeed, I will have ensured its appearance by use of the quasi-performative utterance "I'm trusting you to...". And the second condition required is that in revealing my reliance in this manner I must be expecting that it will engage the dispositions of my trustees, giving them an extra motive or reason for being or acting as I am relying on them to be or act.³

² In putting this argument, I am not wanting to legislate for the use of the word "trust". I am perfectly happy to acknowledge that the characterization I provide of interactively dynamic trust does not catch every variation in usage, even in the usage of the word beyond the limit where it clearly means little more than 'rely'. My primary interest is in demarcating a phenomenon that is clearly of particular interest in human life.

³ Providing an extra motive or reason, as discussed in Pettit 1995, need not mean making it more likely that the person will behave in the manner required; he or she may already have

I think that trust involves dynamic reliance of this kind, because the dynamic aspect provides a nice explanation for why trusting people involves treating them as trustworthy. If I am said to treat you as trustworthy, then I must be treating you in a way that manifests to you—and to any informed witnesses—that I am relying on you; otherwise it would not have the gestalt of treating you as trustworthy. And if I am said to treat you as trustworthy then, in addition, I must be manifesting the expectation that this will increase your reason for acting as I rely on you to act. The implicature of anyone's saying that I treat you as trustworthy is that I expect you to live up to the trust I am investing in you: that is, I expect that the fact that I am relying on you—the fact that I am investing trust in you—will give you more reason than you previously had to display the trait or behaviour required.

Relying on others is a generic kind of activity, then, trust in the sense in which I am concerned with it is a specific form of that generic kind. The difference that marks off trust from reliance, contrary to the suggestion mentioned, is not that trust is a non-rational version of reliance. Rather it is that trust is interactively dynamic in character. It materializes so far as the act of reliance involved is presented as an act of reliance to the trustee, and is presented in the manifest expectation that that will give the trustee extra reason to conform to the expectation.




Figure 1

What of the connection to rationality? I argue that both reliance in general, and trust in particular, may be rational or irrational. While we illustrated reliance on other people by instances that were intuitively rational in character,

enough reason and motive to ensure that they will behave in that way. I can raise the utility that a certain choice has for you, even when it already has much greater utility than any alternative.

nothing rules out cases of irrational reliance. Reliance will be irrational so far as the beliefs on which it is based are not well grounded or, perhaps a less likely possibility, so far as they do not provide a good ground for the reliance that they prompt. Trust is the species of reliance on other people that is interactively dynamic in the sense explained and while this may certainly be irrational, it should be clear that it may be rational too. There may be good reason to expect that others will be motivated by my act of manifest reliance on them and so good reason to indulge in such reliance. I may think that they are not currently, independently disposed to act as I want them to act, for example, but believe that my revealing that I am relying on their acting in that way will provide them with the required motive to do so.

The results we have covered are summed up in the tree diagram attached and I hope that this will help to keep them in mind.

3. Two Forms of Trust

There are two broadly contrasted sorts of beliefs on the basis of which you might be led to trust others in a certain manner: say, trust them to tell the truth on some question, or to keep a promise, or to respond to a request for help.

You might believe that certain others are indeed trustworthy, in the sense of being antecedently disposed to respond to certain manifestations of reliance. They may not be disposed antecedently to display the trait or behaviour you want them to display but they are disposed to do so, other things being equal, should you manifestly rely on them to do so. They are possessed of stable, ground-level dispositions that you are able to engage by acts of manifest reliance.

The dispositions in which you believe in this way would constitute what we normally think of as virtues. You might believe that it is possible to engage some people by manifesting reliance on them because they are loyal friends or associates, for example; or because they are kind and virtuous types who won't generally want to let down someone who depends on them; or because they are prudent and perceptive individuals who will see the long-term benefits available to each of you from cooperation and will be prepared, therefore, to build on the opportunity you provide by manifesting reliance on them.

But there is also a quite different sort of belief that might prompt you to trust certain others, manifesting reliance in the manifest expectation that they will prove reliable. You might think, not that those people are currently disposed to respond appropriately, but rather that they are disposed to form such a disposition under the stimulus provided by your making a relevant overture of trust. You might think that they are meta-disposed in this fashion to tell you the truth, or to keep a promise you elicit, or to provide some help you request. They may not be currently disposed in such directions but they are disposed to become disposed to respond in those ways, should you make the required overture.

This second possibility is less straightforward than the first and I will devote the rest of this section to elucidating it. The possibility is not just a logical

possibility that is unlikely to be realized in practice. It materializes in common interactions as a result of people's desiring the good opinion of others and recognizing—as a matter of shared awareness, however tacit—that this is something that they each desire. It has a salient place within what Geoffrey Brennan and I describe as the economy of esteem (Brennan/Pettit 2004a).

There are two fundamentally different sorts of goods that human beings seek for themselves. The one kind may be described as attitude-dependent, the other as action-dependent (Pettit 1993, chapter 5). Attitude-dependent goods are those which a person can enjoy only so far as they are the object of certain attitudes, in particular certain positive attitudes, on the part of others, or indeed themselves. They are goods like being loved, being liked, being acknowledged, being respected, being admired, and so on. Action-dependent goods are those which a person can procure without having to rely on the presence of any particular attitudes in themselves or others; they are attained by their own efforts, or the efforts of others, and they are attained regardless of the attitudes at the origin of those efforts.

Action-dependent goods are illustrated by the regular sorts of services and commodities and resources to which economists give centre stage. But it should be clear that people care also about goods in the attitude-dependent category; they care about being cherished by others, for example, and about being well regarded by them. Thus Adam Smith, the founding father of economics, thought that the desire for the good opinion or esteem of others, the desire for standing in the eyes of others, was one of the most basic of human inclinations.

“Nature, when she formed man for society, endowed him with an original desire to please, and an original aversion to offend his brethren.

She taught him to feel pleasure in their favourable, and pain in their unfavourable regard. She rendered their approbation most flattering and most agreeable to him for its own sake; and their disapprobation most mortifying and most offensive.” (Smith 1982, 116)

In arguing that people care about the esteem of others, Smith was part of a tradition going back to ancient sources, and a tradition that was particularly powerful in the seventeenth and eighteenth centuries (Lovejoy 1961). I am going to assume that he is right in thinking that people do seek the good opinion of others, even if this desire is not any more basic than their desire for material goods. More particularly, I am going to assume that while the good opinion of others is certainly instrumental in procuring action-dependent goods, it is not desired just as a currency that can be cashed out in action-dependent terms. People will naturally be prepared to trade off esteem for other goods but the less esteem they have, the more reluctant they will be to trade; esteem is an independently attractive good by their lights, not just a proxy for material goods.⁴

⁴ It may be that esteem is desired intrinsically, as a result of our evolutionary priming; or it may be that it is desired instrumentally, where one of the goods it instrumentally promotes is the non-material good of facilitating self-esteem. I abstract from such issues here (see Brennan/Pettit 2004, chapter 1).

The desire for esteem can serve in the role of the meta-disposition of which we spoke earlier. Let people want the esteem of others and they will be disposed to become disposed to prove reliable in response to the trusting manifestation of reliance. Or at least that will be the case in the event that the trusting manifestation of reliance normally serves to communicate a good opinion of the trustee. And all the evidence suggests that it does serve this purpose, constituting a token of the trustor's esteem.

The act of relying on others in a suitable context is a way of displaying a belief that they are not the sort to let you down: they are trustworthy, say in the modality of loyalty or virtue or prudence/perception. The trustor does not typically utter words to the effect that the trustees are people who will not let the needy down: that the trustees, as we say, are indeed trustworthy individuals. But what the trustor does in manifesting reliance is tantamount to saying something of that sort. Let the context be one where, by common assumption, the trustor will expect the trustees to prove reliable in a certain way only if they have a modicum of trustworthiness: only if the trustees are loyal or virtuous or prudent/perceptive or whatever. Under such a routine assumption—more below on why it is routine—the act of trust will be a way of saying that the trustees are indeed of that trustworthy sort.

Indeed, since words are cheap and actions dear, the act of trust will be something of even greater communicative significance. It will communicate in the most credible currency available to human beings—in the gold standard represented by action—that the trustor believes the trustees to be truly trustworthy: to be truly the sorts of people who will not take advantage of someone who puts himself or herself in their hands. It does not just record the reality of that belief, it shows that the belief exists. Thus Hobbes (1991, 64) can write: "To believe, to trust, to rely on another, is to Honour him: sign of opinion of his vertue and power. To distrust, or not believe, is to Dishonour."

When it connects in this way with the desire of a good opinion, the act of trust has an important motivating aspect for the trustees. It makes clear to them that they enjoy the good opinion of the trustor—the belief that they are trustworthy—but that they will lose that good opinion if they let the trustor down. This means that the trustor has a reason to expect the manifestation of reliance to be persuasive with the trustees, independently of any belief in their pre-existing loyalty or virtue or prudence. If the trustees value the good opinion of the trustor, which the manifestation of reliance reveals, then that is likely to give them pause about letting the trustor down, even if they are actually not particularly loyal or virtuous or prudent/perceptive in character. Let the trustor down and they may gain some immediate advantage or save themselves some immediate cost. But let the trustor down and they will forfeit another immediate advantage: the salient benefit of being well regarded by the trustor, as well as the other benefits associated with enjoying such a status.

But there is also more to say. When I display trust in certain others, I often demonstrate to third parties that I trust these people. Other things being equal, then, such a demonstration will serve to win a good opinion for the trustees among those witnesses: the demonstration will amount to testimony

that the trustees are trustworthy. Indeed if the fact of such universal testimony is salient to all, the demonstration may not just cause everyone to think well of the trustees; it may cause this to become a matter of common belief, with everyone believing it, everyone believing that everyone believes it, and so on. Assuming that such facts are going to be visible to any perceptive trustees, then, the existence of independent witnesses to the act of trust will provide further esteem-centred motives for them to perform as expected. Let the trustor down and not only will trustees lose the good opinion that the trustor has displayed; they will also lose the good opinion and the high status that the trustor may have won for them among third parties.

The belief that someone is loyal or virtuous or prudent/perceptive may explain why the risk-taking that trust involves may actually be quite sensible or rational. Certainly there is a risk involved in this or that act of trust but the risk is not substantial—it is, at the least, a rational gamble—given that the trustee has those qualities.⁵ What we now see is that the belief that certain parties deserve esteem, and that responding appropriately to an overture of trust will secure esteem for them, may equally explain why it is rational to trust those people.⁶ It does not direct us to any independent reason why the trustees may be taken to be autocratically reliable—any reason of objective trustworthiness—but it reveals how the act of trust can transform the trustee into reliable parties, eliciting the disposition to perform appropriately. To manifest trusting reliance is to provide normal, esteem-sensitive trustees with an incentive to do the very thing which the trustor is relying on them to do. It is a sort of bootstraps operation, wherein the trustor takes a risk and, by the very fact of taking that risk, shifts the odds in his or her own favour.

Believing that certain individuals are loyal or virtuous or prudent/perceptive is quite consistent, we should notice, with believing that still in some measure they desire the esteem of others. This is important because it means that people may have double reason for trusting others. They may trust them both because they think that they are trustworthy and—a back-up consideration, as it were—because they think that they will savour the esteem that goes with proving reliable and being thought to be trustworthy. I said earlier that to trust certain others is to treat them as trustworthy. When one trusts them in the standard way, one treats them as trustworthy in the sense of acting out of a belief that they are trustworthy. When one trusts them on the esteem-related basis, one

⁵ For the record, I think that the risk involved in the act of trust need not be a risk of the ordinary, probabilistic kind (Pettit 1995). Take the case where I am dealing with others whom I believe to be more or less certain of responding appropriately to an act of manifest reliance on my part; let my degree of confidence that they are reliable in this way be as near as you like to 1. I can still be said to trust such people, so far as I put my fate in their hands when I rely on them. I do not expose myself to a significant probability that they will betray me—that probability may approach 0—but I do expose myself to the accessibility of betrayal to them; I expose myself to their having the freedom to betray me. Here I break with Russell Hardin 1992, 507, and join with Richard Holton 1994.

⁶ I have come to realize, from discussions with Victoria McGeer, that the role of belief here may be played by the attitude of hope, as I have characterized it elsewhere (Pettit 2004). For an exploration of this idea see McGeer 2004b.

treats them as if they were trustworthy, whether as a matter of fact they are trustworthy or not.

One final issue. The esteem-related way in which trust may materialize depends on its going without saying—its being a matter of routine assumption shared among people—that when a trustor invests trust in a trustee, that is because of taking the trustee to be trustworthy. But isn't it likely that people will recognize that in many cases the trustor invests trust because of taking the trustee to want his or her esteem, or the esteem of witnesses, not because of taking the person to be antecedently trustworthy? And in that case won't the mechanism we have been describing be undermined? People are not going to expect to attract esteem for proving reliable, if they expect that their proving reliable will be explained by the trustor, and by witnesses, as an effect of their wanting to win that esteem. They will expect to attract esteem only if they think that their proving reliable will be generally explained by the assumption that they are trustworthy types: by the attribution of stable dispositions like loyalty or virtue or prudence/perception.

Is there any special reason to think that the system won't unravel in this way, and that it will continue to go without saying—it will continue to be a matter of common assumption—that people who prove reliable under overtures of trust will enjoy the attribution of estimable, trustworthy dispositions? I believe there is. The assumption is going to remain in place as long as people are subject to the fundamental attribution error or bias, as psychologists call it, and so are likely to expect everyone to conform to that pattern of attribution. And a firm tradition of psychological thought suggests that the bias is deeply and undeniably ingrained in our nature.

E. E. Jones (1990, 138) gives forceful expression to the view that the bias has this sort of hold upon us: "I have a candidate for the most robust and repeatable finding in social psychology: the tendency to see behavior as caused by a stable personal disposition of the actor when it can be just as easily explained as a natural response to more than adequate situational pressure." This finding—that people are deeply prone to the fundamental attribution bias—supports the idea that, even if they are conscious of their own sensitivity to a force like the desire for esteem (Miller/Prentice 1996, 804), people will be loathe to trace the behaviour of others to such a situational pressure. They are much more likely to explain the behaviour by ascribing a corresponding disposition to them. And that being so, they are likely to expect each to do the same, to expect that each will expect each to do the same, and so on in the usual hierarchy. Thus they are likely to expect that trustors will invest trust in certain others only so far as they take those others to have the stable personal dispositions associated with trustworthiness.⁷

⁷ A related problem arises with the trustor as distinct from the trustee. Why should the trustor expect that the trustee and other witnesses will take them to be moved, not by a wish to signal esteem and thereby motivate the trustee, but rather by the attribution of a trustworthy disposition to the trustee? The answer, I think, is that to the extent that people tend to explain the esteem-seeking behaviour of others by attributing stable dispositions they will also tend to explain the relevant sort of esteem-signalling behaviour as springing from the attribution of such dispositions. They will display, not just an attribution bias, but a meta-

4. The Internet

And so, finally, to the connection between trust and the Internet. The question that I want to raise is whether the Internet offers a milieu within which relations of trust—trust as distinct from reliance—can rationally develop. There is every reason, of course, why people who already enjoy such relations with one another should be able to express and elicit trust in one another over the Internet. But the question is whether the Internet offers the sort of ecology within which trust can rationally form and strengthen in the absence of face-to-face or other contact. Is it a space in which I might rationally make myself reliant on others by sharing difficult secrets, asking their advice about personal problems, exposing myself financially in some proposal, and so on?

We distinguished in the last section between two sorts of bases on which trust may emerge in general. The primary basis for trust is the belief that certain people are trustworthy: that is, have stable dispositions like loyalty and virtue and prudence/perception. Primary trust will be rational just in case that belief is rational and serves rationally to control what the trustor does. The secondary basis for trust is the belief that even if the people in question are not trustworthy—even if they do not have stable dispositions of the kind mentioned—they are meta-disposed to display the trait or behaviour that the trustor relies on them, now in this instance, now in that, to display. More concretely, they desire esteem and they can be moved by the esteem communicated by an act of trust—and perhaps broadcast to others—into becoming disposed to be or act as the trustor wants them to be or act. The secondary form of trust that is prompted in this manner will be rational just in case the belief in the esteem-seeking meta-disposition is rational and serves rationally to shape the trustor's overture.

Does the Internet offer a framework for the rational formation of primary trust? In particular, does it provide an environment where I may rationally come to think that someone I encounter only in that milieu is a likely to respond as a loyal or virtuous or even prudent/perceptive person? Or does it offer a framework for the rational formation of secondary trust? Does it enable me to recognize and activate another's desire for esteem, creating a ground for expecting that he or she will respond favourably to my trusting displays of esteem?

There is no problem with the possibility of the Internet facilitating rational reliance, as distinct from trust. Suppose I become aware of someone over email or in a chat room or via the web. And imagine that an opportunity arises where I will find it rational to do something—say, go to a proposed meeting place—only if there is reason to believe that the other person will act in a certain way: in this case, be at the proposed place to meet me. I may not have very much solid evidence available about that person over the Internet—deception is not easily detectable—but there is nothing to block the possibility that what evidence I have makes it rational for me to rely on their doing this or that; what evidence I have makes that a rational gamble.

attribution bias: a tendency to take people to employ an attributionist heuristic in interpreting and dealing with others.

But reliance is one thing, trust another. Take the question of primary trust first of all. Is it ever likely to be the case, with the individuals I encounter on the Internet, and on the Internet only, that I can come to think of them as loyal or virtuous or even prudent/perceptive: that is, capable of recognizing and responding to a sense of the long-term interests that they and I may have in cooperating with one another? And is it ever likely to be possible for me to invest trust rationally in such contacts?

I think not. Consider the ways in which I come to form beliefs about loyalty and virtue and prudence/perception in everyday life. I may rely in the formation of such beliefs on at least three distinct sources of evidence. First, the evidence available to me as I see and get cues—no doubt at subpersonal as well as personal levels of awareness—to the expressions, the gestures, the words, the looks of people: in a phrase, their bodily presence. Call this the evidence of face. Second, the evidence available to me as I see the person in interaction with others, enjoying the testimony of their association and support: in particular, see them connecting in this way with others whom I already know and credit. Call this the evidence of frame. And third, the evidence that accumulates in the record that I will normally maintain, however unconsciously, about their behaviour towards me and towards others over time. Call this the evidence registered in a personal file on the people involved.

The striking thing about Internet contact is that it does not allow me to avail myself of such bodies of evidence, whether of face, frame or file. The contact whose address and words reach my screen is only a virtual presence, albeit a presence I may dress up in the images that faustasy supplies. I cannot read the face of such a contact; the person is a spectral, not a bodily presence in my life. Nor can I see evidence of his or her character—and I won't be able to establish independently whether 'his' or 'her' is appropriate—in the interactions the person enjoys with other persons familiar to me, assuming that such witnesses will be themselves only spectral presences in my experience. And nor, finally, will I be able to keep a file on the performance of the person over time, whether with me or with others. There won't be any way of tracking that person for sure, since a given person may assume many addresses and the address of one person can be mimicked by others.

Not only do these problems stand in the way of my being able to judge that a pure Internet contact is loyal or virtuous or prudent/perceptive. They are compounded by the fact that such problems, as I am in a position to see, will also stand in the way of others when they try to read and relate to me. For them I will be just a spectral presence, as they are spectral presences for me. Our voices may call out over the Internet, but it won't ever be clear where they come from or to whom they belong. They will be like a chorus of lost cries, seeking in vain to pin one another down. Or at least that is what they will be like, absent the illusions that fantasy may weave as it claims to find structure and stability in the shifts of the kaleidoscope.

On the Internet, to put these problems in summary form, we all wear the ring of Gyges. Plato took up an old myth in asking whether we would be likely to remain virtuous, did we have access to a ring that would give us power, on

wearing it, to become invisible and undetectable to others. That myth becomes reality on the Internet for, with a little ingenuity, any one of us may make contact with another under one address and then, slipping that name, present ourselves under a second or third address and try to manipulate the other's responses to the first. That we exist under the second or third address may not be undetectable to the other in such a case but that it is we who do so—that we have the same identity—certainly will be undetectable.

In view of these difficulties, I think that the possibility of rational, primary trust in the virtual space of the Internet is only of vanishing significance. It is a space in which voices sound out of the dark, echoing to us in a void where it is never clear who is who and what is what. Or at least that is so when we enter the Internet without connection to existing, real-world networks of association and friendship.

But what of secondary trust? Are the prospects any better here that we will be able to reach out to one another in the environment of the Internet and forge relationships of trust? I think not. I may be able to assume, as a general truth about human nature, that those with whom I make contact are likely to savour esteem, including the esteem of someone like me that they don't yet fully know. But how can I think that anything I do in manifesting reliance will seem to make esteem available to them, whether it be my own esteem or the esteem of independent witnesses?

The problem here derives from the problems that jeopardize the possibility of primary trust. I am blocked from rationally forming the belief that an Internet contact is loyal or virtuous or prudent/perceptive, as we saw. Since that blocking is something that anyone is in a position to recognize, others will see that it is in place, and I will be positioned to recognize that they will see this. And that being so, I will have no ground to think that others—other pure Internet contacts—are likely to take an act of manifest reliance on my part as an expression of the belief that they are people of proven or even rationally presumed loyalty or virtue or prudence/perception. I will have no ground for expecting them to take my act of trust as a token of personal esteem.

Nor is there any way out of this difficulty available by recourse to the possibility of witnesses. For as those on whom I bestow trust will be unable to see my manifestation of reliance as a token of esteem, so the witnesses to my act will be equally prevented from viewing it in that way. The addressees and the witnesses of my act may see it as a rather stupid, perhaps even pathetic attempt to make contact. Or if the act promises a possible benefit to me at a loss to the addressees—as in the case of emails that propose various financial scams—they may see it as a rather obvious attempt at manipulation and fraud. And that just about exhausts the possibilities. If I try to invest trust in others unknown to me outside the Internet, then the profile I assume will have to be that of the idiot or the trickster. Not a happy choice.

The claims I have been making may strike some as exaggerated. But if they do, that may be because of a confusion between what I described at the beginning of the paper as Internet trust between real people—my topic here—and Internet trust between Internet people: that is, between personas that we

real individuals construct on Internet forums. If I construct an agony aunt persona on an Internet forum, then in that persona I may succeed over time in earning—earning, not just winning—the trust of those who, in the guise of other Internet identities, seek my guidance. This form of trust is of great interest and opens up possibilities of productive human relationships but it is not the phenomenon that I have been discussing here. My concern has been with how far real people can manage, on the basis of pure Internet contact, to establish trust in one another. And the answer to which I am driven is that they cannot effectively do so. The message of the paper, in a word used by Hubert Dreyfus (2001), is that ‘telepresence’ is not enough on its own to mediate the appearance of rational trust between real people.

One concluding word of caution, however. I have argued for this claim on the assumption that telepresence will remain as Gygean as it currently is: that it will continue to lack the facial salience, the framed support, and the fileable identities available in regular encounters with other people. I am no futurist, however, and I cannot say that telepresence will always remain constrained in these ways. Perhaps lurking out there in the future of our species is an arrangement under which telepresence can assume firmer, more assuring forms and can serve to mediate rational trust. I do not say that such a brave new world is logically impossible. I only say that it has not yet arrived.⁸

Bibliography

- Brennan, G./P. Pettit (2004a), *The Economy of Esteem: An Essay on Civil and Political Society*, Oxford
- / — (2004b), Esteem, Identifiability, and the Internet, in: *Analyse & Kritik*, this volume
- Dreyfus (2001), *On the Internet*, London
- Hardin, R. (1992), The Street-Level Epistemology of Trust, in: *Analyse & Kritik* 14, 152–176
- Hobbes, T. (1991), *Leviathan*, Cambridge
- Holton, R. (1994), Deciding to Trust, Coming to Believe, in: *Australasian Journal of Philosophy* 72, 63–76
- Jones, E. E. (1990), *Interpersonal Perception*, New York
- Lovejoy, A. O. (1961), *Reflections on Human Nature*, Baltimore
- McGeer, V. (2004a), Developing Trust on the Internet, in: *Analyse & Kritik*, this volume
- (2004b), *Trust, Hope and Empowerment*, Princeton
- Miller, D. T./D. A. Prentice (1996), The Construction of Social Norms and Standards, in: E. T. Higgins/A. W. Kruglanski (eds.), *Social Psychology: Handbook of Basic Principles*, New York, 799–829
- Pettit, P. (1993), *The Common Mind: An Essay on Psychology, Society and Politics*, paperback edition 1996, New York
- (1995), The Cunning of Trust, in: *Philosophy and Public Affairs* 24, 202–25

⁸ My thanks to Geoffrey Brennan and Victoria McGeer for background discussions, to those attending the Bielefeld conference for helpful comments, and to Michael Baumann for a very insightful set of queries and suggestions.

- (2004), Hope and its Place in Mind, in: *Annals of the American Academy of Political and Social Science*, (forthcoming)
- Smith, A. (1982), *The Theory of the Moral Sentiments*, Indianapolis