

Russell Hardin

Internet Capital*

Abstract: The Internet is a huge form of social capital that is not reducible in its characteristics to other forms of social capital, such as ordinary networks of people who more or less know each other. It enables us to do many things with radically greater efficiency than we could without it. It can do some things better but other things much less well than traditional devices can. At both extremes, the differences are so great as to be not merely quantitative but also qualitative. The things it can do better include things that can readily be checked and verified. The things that it often cannot do include securing commitments for action. A brief history of the forms of social cooperation suggests that relationships on the Internet are typically too thin to back trust and cooperation among those who do not have fairly rich relationships off-line.

0. Introduction

The Internet is a form of social capital that is not reducible in its characteristics to other forms of social capital, such as ordinary networks of people who more or less know each other, organizations whose connections and talents may be tapped by ‘users’, and so forth. The social capital of the Internet—which I will call Internet capital to distinguish it from network and organizational capital—is huge. It enables us to do many things with radically greater efficiency than we could do them without the Internet as an enabler. Because the Internet itself is enormous, with millions of participants in it, a typical individual’s contribution to its capital is negligible. Many other kinds of network are far smaller with far fewer participants. In these, an individual’s contribution can proportionately be relatively large, so that my using the network might increase its value to all other users. My contributions to the wealth of resources available on the web, however, is proportionately tiny. There are some atypical participants who make major contributions that are enhanced by the very nature of the Internet in its capacity to reach millions of people at virtually no marginal cost to the person providing a benefit. If you program a fix to, say, one of the notoriously frequently discovered flaws of any of the Microsoft operating systems, you can make that fix available to millions of people. That fix is virtually a pure public

* I thank Nomy Arpaly and Azi Lev-On for discussions of these issues, and I thank Michael Baurmann for an extensive written commentary and for several recommendations to expand arguments. I also thank Mariel Ettinger and Huan Wang for creative and energetic research assistance under the duress of a ticking clock.

good whose marginal cost to an additional user is zero and whose benefit could be great for every one of millions of users.

In an informal conversation many years ago, Thomas Schelling estimated that the computer as word processor had more or less doubled the productivity of the more productive academics (doubling the productivity of some would add nothing to their total). One might now ask how great has been the effect of the Internet as a research tool, at least in certain fields, and perhaps in all academic and journalistic fields. My experience is perhaps indicative, even though I am not especially adept. I can now check vague memories of what someone wrote in a matter of minutes. I have a younger colleague who checks the Internet while we are talking by telephone and she finds sources for claims and the accurate wording of the claims almost instantly. Most of my generation are not that adept, but even we can do what formerly would have been amazing things. In some areas it may prove to have enormous impact. For example, it enables lawyers on both sides of a case to find scores of prior decisions that fit the decisions they want in the case. Judge Richard Posner (private conversation) speculates that this will undermine the principle of *stare decisis* in the common law.

For all its magic, however, the Internet has what seem to be major drawbacks that make it not merely a high-speed analog of traditional devices. It can do some things better but other things much less well than traditional devices can. At both extremes, the differences are so great as to be not merely quantitative but also qualitative. The things it can do better are things that can readily be checked and verified. The things that it often cannot do include securing commitments for action, whether for spontaneous cooperation or well organized cooperation. To see both the limits and the possibilities, consider a brief social history of cooperation.

2. A Social History of Cooperation

In pre-historical times, people evidently lived in small groups of a hundred or two hundred people, mostly in extended families. In such a context, social cooperation can be readily organized by the tightly interwoven interests of everyone in the group. Each of us can enforce behavior according to communal norms. We benefit from the monitoring of all by all and probably also from the badgering and hectoring of all miscreants by everyone else. Because each of us shares many interests with each other person in the group, we have strong incentive to be cooperative in order to maintain the cooperative inclinations of the others. But the tendency we know in modern times to generalize from such interests to some overall sense of the rightness of relevant behaviors might have affected even very primitive groups, which might have moralized behaviors that serve interests. A small society, and not only a primitive small society, might not distinguish between normatively and pragmatically correct rules for behavior. Free-riding will be subject to sanction in a very small, close society, so that cooperative action can be over-determined by the combination of moral and self-interest motivations; your moral commitment to the group and its benefits from cooperation

will be reinforced by the fact that it will be in your interests to cooperate in order to avoid sanctions. In a larger society, large enough that each individual does not know all the others, or even a large fraction of them, there is likely to be some need for regulation of relationships beyond those whom we know well and with whom we have extensive ongoing relationships. Without such regulation, free-riding can undercut all large-scale cooperative ventures.

This brief history parallels David Hume's account of the evolution of social order. His account is what Rudolf Carnap (1962, 576-7) would call a rational reconstruction. It is supposed to be a plausible story of social evolution because it is based on an articulate explanatory account of how cooperation can be sustained at various scales of society. It starts with two evident facts: that orderly, stable societies have existed at the very small scale and at the very large scale that permits anonymity toward most people. We need explanations of the order at these two scales of society and, for some purposes, we might also want to have an account of the evolution from small to large scale. The latter is not important here; we need only to understand how order can be achieved at these two scales so that we can then determine what kind of order to expect for life on the Internet at its grandiose scale. It is maintenance, not creation, of order that we need to explain in order to grasp how cooperation works in our lives.

On Hume's account, social order can take two quite different forms. First, in a very small society such as that of a small Indian tribe of North America in Hume's time, order can follow from conventions that can govern behavior in conditions of the general transparency of everyone's actions to everyone else. Such a society can be organized by norms, although this is not Hume's term (the term with its modern force would have been unavailable in his time).¹ He speaks of conventions. He is in fact the inventor of the idea of a convention that is the result of repeated coordination on one of the possible coordination outcomes for various interactions, such as the rules for determining the right of succession of monarchs, the rules of property, the rules of the road, and so on. Once we establish such a convention, such as that for driving on the right in North America, that convention is self-reinforcing in the very strong sense that each individual will find it in her interest to abide by or acquiesce in the convention. Many norms are *de facto* coordination conventions (Lewis 1969; Hardin 1982, chapters 10-14; and 1995, chapters 2-4).

In Thomas Hobbes's state of nature there is neither justice nor property (nor right nor wrong). Hume agrees that in a Hobbesian state of nature there would be neither justice nor injustice, there would be no property, and life would be grim. But he thinks this is an "imaginary state", a philosopher's device, not a historical state of affairs (Hume [1739-40] 1978, 3.2.2, 501). For him there is an intermediate position between this state of nature (which did not ever exist) and a large society such as ours with its elaborate government. That intermediate position is a small society that is regulated by conventions. Those conventions might add up to some perhaps minimal sense of justice. This difference in views

¹ The first use of "norm" cited in the *Oxford English Dictionary* is in 1821 and the term is used in the sense of a standard, model, pattern, or type. The next use in 1828 is in the relevant normative sense.

turns on Hume's richer theoretical grasp of strategic possibilities in his theory of convention as applied to the iterated interactions of a small society (Hardin 2004, chapters 1 and 9).

Instead of Hobbes's war of all against all in a state of nature, in a small society we face a substantially contrary condition: *the monitoring of all by all*. And we are therefore all basically reliable in our dealings with each other and toward our collective interests. As we now know from many anthropological studies of small, relatively primitive societies, they can work without any heavy machinery of government and with relative spontaneity of action by everyone. They enjoy spontaneous order.

Hume's second form of social order is that of a large and at least moderately complex society in which we are not all known to each other and in which monitoring and individual-level sanctioning of miscreants cannot be expected to guarantee generally good behavior by all. It would be impossible to regulate such a society with nothing more than norms and reciprocity. In such a society we need government to restrain us from harmful actions and probably even to compel us to beneficial collective actions.

In our much larger society, we too can enjoy spontaneous order in many relationships, especially those that are ongoing. But if our exchange relations are restricted to the small numbers with whom we can repeatedly interact, we face a serious loss of opportunities that we could enjoy if we could also guarantee reciprocal fulfillment of even isolated exchanges. Moreover, even in ongoing relationships, we cannot trust one another to abide by exchanges that involve very large values, so that our relationships will still be restricted. For example, to whom would you sell your house on a legally unenforceable contract to pay you a large monthly sum for the next twenty years? Hence, even if we do not go all the way with Hobbes in thinking that unregulated social interactions would be constantly murderous, as in his state of nature, we must agree with him that they would be radically poorer than what we could have under a properly functioning government.

In both these states there must be some sense of justice that provides for social order. Under a regime of such justice, there could arise principles of possession, although in the conditions of the primitive tribe the range of things that could come under the rubric of possession might be extremely small and restricted. In a nomadic pastoral society, our goats, sheep, or cattle might count as individual possessions. In a hunter-gatherer society, almost nothing might count as individual possessions. In neither case would land count as a possession, although my tent might describe an area onto which you should not encroach while we are in this particular place.

Note that the small primitive society could lack government and could be relatively anarchic but nevertheless well-ordered, with each of us free to do as we please much of the time but with very clear expectations about behaviors in many contexts of importance to our group. As anarchists have long insisted, there can be order in anarchy. Hume says, "An *Indian* is but little tempted to dispossess another of his hut, or to steal his bow, as being already provided of the same advantages; and as to any superior fortune, which may attend one above

another in hunting and fishing, 'tis only casual and temporary, and will have but small tendency to disturb society" (Hume [1739–40] 1978, 3.2.8, 539). Hence, despite its lack of government, the society need not descend into the violence and virtual war of Hobbes's state of nature. I have nothing to gain from attacking or stealing from my neighbor. Indeed, if I do steal, my theft would commonly be known to all in the society and I might suffer powerful sanctions.

Hume therefore specifically says that he disagrees so much with the Hobbesian vision of the state of nature that he supposes that government arose not in order to deal with disorder within society but to allow for better organization of defense against those outside the society who might attack it (Hume [1739–40] 1978, 3.2.8, 539–40). Hume says the "state of society without government is one of the most natural states of men, and may subsist with the conjunction of many families, and long after the first generation" (*ibid.*, 541). The difference between Hobbes and Hume here might primarily be in the greater body of empirical evidence on the North American Indians when Hume wrote, almost exactly a century after Hobbes first wrote.² Hobbes presumably had no evidence for his speculative claim that the Indians lived in the brutish manner of a war of all against all (Hobbes [1651] 1968, chapter 13, [63] 187).

We now have extensive evidence on societies that manage without powerful governments. We also have the evidence Hobbes had that within our highly governed societies, we can have unregulated exchanges without such fear of being abused as to dissuade us from attempting to cooperate with others. Hume seems to understand both of these arenas for successful cooperation. On the second, his evidence is not more impressive than Hobbes's. Therefore, his better understanding seems to come from better theory, namely his theory of convention.

Although it would be possible to maintain a very small society indefinitely without government, Hume ([1739–40] 1978, 3.2.8, 541) says, "tis impossible they shou'd maintain a society of any kind without justice", which entails three fundamental principles concerning (1) the stability of possession, (2) its transfer by consent from one owner to another, and (3) the performance of promises (or contracts loosely defined).³ Hume calls these laws of nature because they are *de facto* antecedent to government in the sense that they are sociological principles. Hume's three laws of nature are elevated to their high status by the fact that they are enormously useful in enabling us to have stable expectations and to invest our efforts in individually beneficial endeavors that have beneficial implications for the whole society. Seen *ex ante*, his three laws are mutually beneficial. Adherence to each of these laws is a *sine qua non* of cooperative social order.

We can fully understand these laws of nature and still act as though they did not govern our behavior. Indeed, in all three cases, it is not even generally

² We now have extensive evidence on various forms of societal development. See, for example, Johnson/Earle 2002.

³ Promising is more important here than it might seem because it covers contracting as well and it includes a large fraction of exchanges, many of which cannot be made on the spot but must be negotiated for some future time, especially if what is promised by at least one of the parties is performance of some action or plan, such as building a house or some lesser object.

clear what it would mean for a single individual to act according to the derived laws or principles.⁴ Hume says, “I should be the cully of my integrity, if I alone shou’d impose on myself a severe restraint amidst the licentiousness of others” (Hume [1739-40] 1978, 3.2.7, 535). That would be too much to expect, especially of a rational, moderately self-interested person. One can be obliged to follow the dictates of the laws of nature in only two contexts: first in a society small enough that it can be regulated by spontaneously enforced and obeyed norms; second in the context of a political society under a government that will enforce them as legal laws.

Hume supposes that government would get its first authority from these principles, which could be established before formal government arises (T3.2.8, 541). It is only with greater prosperity and larger scale that we might come into such serious conflict with each other as to need government to regulate our behavior. Although adherence to the principles of justice would be sufficient to maintain any society, it would be impossible for us to observe those principles in a large society unless we have government to enforce compliance with the principles (543). Hume’s general concern is with social order, for which stable property relations are necessary but, of course, not sufficient. Moreover, in addition to his ‘laws’ of property, which we might call norms, there could be many other norms to regulate behavior and to help secure social order. Such norms might lie outside the law while supporting law. For example, our small community might have strong norms regulating cooperation in many contexts and our norms might differ from those of a neighboring community that is under the same government and laws.

3. Trust as Encapsulated Interest

Trust evidently arose historically in response to a world in which we do not know everyone but need nevertheless to rely on many others for various things. The very term trust is a late addition to the English language and it has yet to gain full currency—with both a noun and a verb form—in many languages (Hardin 2002, 57-8). Substantial division of labor in essence created the need for making sure that others will be reliable while they do things that are complementary to what one does oneself. In many contexts in modern lives, we cannot even judge the competence of those who do particular and very important things for us. For example, most of us cannot fully judge the competence of our doctors, lawyers, or other professionals, or of agents of various organizations, especially including government. Yet we need them because we do not know what they know and often cannot judge the validity of what they do. The students at the universities at which I have taught often cannot generally judge the competence of their

⁴ Hobbes goes so far as to say that it would be wrong for a single person to follow the laws of nature in the actual world if no one else were following them because to do so would make one a prey to others and procure one’s own certain ruin, “contrary to the ground of all Lawes of Nature, which tend to Natures preservation” (Hobbes [1651] 1968, chapter 15, [79] 215).

teachers. What they often can judge is how moved they are by the teachers, or how entertained.

There are many accounts of trust that make it quite varied in content, possibly more varied in academic treatments than even in the vernacular. For the issues here, the most useful conception is of trust as encapsulated interest, as follows. If you and I have an ongoing relationship in which we reciprocally do things for each other, and if I am quite sure that you value that relationship enough to want to continue it, then I trust you because you *de facto* encapsulate my interests in your own. That does not mean you will never violate my trust or even that I would believe you will never violate it. It merely means you have a strong interest in fulfilling my trust (Hardin 2002). Clearly, we cannot be in vast numbers of relationships such as this. Hence, trust as encapsulated interest cannot regulate our interactions with everyone. I will henceforth mean only this form of trust.

Some of the other accounts of trust are utterly implausible as explanations of why we behave cooperatively. Most of them inadvertently are accounts of trustworthiness, even though they are wrongly labeled as theories of trust. Among credible alternative accounts are trust that is grounded in an assessment of the other person's moral or psychological commitments. If I know you are morally committed to fulfilling any agreement for cooperation that you enter into either explicitly or implicitly, then I can trust you in this sense. If I think you have a strong commitment to being the kind of person who is reliable in cooperative ventures, similarly, I can trust you. Neither of these conceptions of trust can likely apply to many, if any, relationships over the Internet or even relationships offline. Trust as encapsulated interest is far more common than either. But each of these three conceptions probably applies to some interactions.⁵

In general, on the Internet, trust as encapsulated interest is likely to work only if we have an offline relationship that gives us each more knowledge about each other than we could typically get from merely an Internet relationship (without a nagging worry whether all that knowledge is true). Our problem here, therefore, is to determine how we can come to be confident of our Internet relationships. We also face that problem for our dealings with many, perhaps most people offline. We can canvass the devices that help us offline to see whether they could also work online. Many of them cannot, because our relationships on the Internet are too focused and not thick enough.

One of the devices we have to secure trustworthy behavior from those on whom we have to rely, even though we would not trust them in any strong sense, is to have third-party certification of some kind. This can range from the advice of a trusted neighbor who says that a plumber is good to the backing of a professional or institutional body that tests or has opportunity to evaluate competence. It could also include cases in which the third party more or less guarantees the performance of someone we are in no position to trust but with whom we would like to be able to cooperate on something.

⁵ The trustworthiness-as-character and trustworthiness-as-moral-commitment theories cannot be applied to the person whom you trust and I distrust unless one of us is mistaken. Clearly, we can often have such contrary relationships with another.

It is commonly noted that trust arises where there is the substantial risk of default, or untrustworthy behavior. If there were exceedingly little risk that another would default—for example, if the other were acting at gunpoint to do our bidding—it would be pointless to speak of trust. If we are to establish trust in some relationship, then we will have to face some risk, repeatedly, that means we have a real stake in the judgment of the relationship. With no risk, we might wrongly suppose we come to know someone well enough to judge how they will behave. But if we have faced them repeatedly in contexts in which we are both at risk, we can finally know with great confidence whether the other is reliable and—perhaps even more important—we can know whether the other has an interest in maintaining our relationship and therefore has not only the competence but also the will to fulfill our trust.

In very anomic contexts, we cannot typically have either kind of information about those with whom we interact: neither the very general knowledge of their competence nor the personally very specific knowledge of their interest in fulfilling our trust.

What are the risks that we take over the Internet that we can successfully handle so that our ‘partners’ on the web do not take unreasonable advantage of us? For commercial dealings, we do take some risks, although we substantially reduce these in the same way we do in commercial dealings off the web. We rely on dealers’ reputations. This is not a backward-looking estimate of their character, but a forward-looking claim about their interests. Reputation is commonly represented as backward-looking. I think this is almost always a mistake for issues that might involve risky exchanges and trust. A store’s reputation is more important by far to the store than it is to me. Therefore the store cultivates a good reputation. The dealers we wish to buy from are those who care about their *future* reputation and who therefore have an interest in dealing equitably with us as a way of investing in their future reputation.

This is the nature of John Mueller’s account of the early history of Wanamaker’s department store in Philadelphia. John Wanamaker saw that he could attract far more customers if he could convince people that his policy was not *caveat emptor*—let the buyer beware—which was the policy of virtually all sellers in his time. He therefore advertised that his store would take back anything it sold for a full refund, no questions asked (Mueller 1999, 79–80). When I deal with Amazon.com, I think of it no differently than I think of a major department store on Lexington Avenue. It is a business that needs repeat customers. And its success suggests that it has figured out how to keep its customers coming back for more.

In my commercial dealings with stores, trust is not an issue. Competence is the whole story. I want them to be competent in serving their customers well. I want them to be competent in understanding that their future reputation is important to them and in understanding that they want repeat customers, not one-time customers whom they could readily cheat without concern for the future. Shopping over the Internet does not add any significant complication over the normal nature of our offline dealings with stores. It is in other areas where trust is absent that the devices for ensuring reliable behavior become

important. As discussed below, *the only major category of such problems might be in the use of the Internet for email for various purposes.*

The most obvious feature of the Internet for this discussion is that it is a vast collection of networks. Offline networks commonly enable us to build up trust relationships. Often these are very limited and focused. Let us therefore specifically address the possibilities for network trust.

4. Trustworthiness in Communities and Networks

Recall Hume's account of the rise of a large society with a government from earlier forms of small communities. In small, close communities cooperation can be managed by the rule of norms that are backed by spontaneous sanctions from the members of the community itself. Hence, those members are essentially trustworthy in their dealings with each other. Even today, many of us live at least parts of our lives in small 'communities' in which norms govern much of our behavior toward each other. But large parts of our lives are not regulated by close communal ties and the spontaneous sanctions to which they give rise. We live in more nearly urban contexts in which we interact with large numbers of people, many of whom we may never expect to see again after a single interaction. More significantly, we have ongoing relationships with extensive networks of people. We secure what we need through multiple networks, some of them overlapping in part but some of them essentially separate from others (see further, Cook/Hardin 2001). In these contexts, we commonly develop trust relationships in the sense of trust as encapsulated interests. We reciprocate favors over many years with people whom we grow to trust because we value our relationships with them and we suppose they value their relationships with us.

Although there are exceptions, most of our relationships over the Internet cannot have the quality of small, close communities that are spontaneously regulated by norms. But they also cannot usually fit the structure of trust relations, because such relations are typically embedded in networks or are built up in dyadic relationships that go on for a long time. It is difficult to imagine devices that would work for trust on the Internet with people who do not have further interactions with each other off the Internet, interactions in which trust could be grounded. We can have smallish networks that have more or less stable membership and that are involved in frequent interactions on the Internet. But the trust relations that normally grow in iterated and network interactions offline grow out of interactions in which there is something at stake, often something very significant at stake. I do things for you and you do things for me and we trust each other with respect to some range of things.

On the Internet, when there is something of significant value at stake, the potential partners to an exchange do not know each other from before and might typically not expect ever to encounter each other again. Only an ongoing, stable institution, such as Amazon.com or the website of a traditional retailer can expect to be judged by its reputation because only such an entity can have enough interactions with others to be able to establish a reputation. Such an

entity does not need to trust its customers, who can obtain goods or services only by first paying for them. Indeed, Internet retailers do not even suffer the steady losses of shoplifting that afflict real stores.

A striking case of intensive activity on the Internet is by those who play Internet games, which are games that are on the web and that players can enter into or exit from more or less at will. Players can number into the millions and can come from anywhere on the globe (Turkle 1996). They are typically anonymous to each other, although each can have a name for purposes of the game so that players know with 'whom' they are playing in the limited sense of knowing whether any player was in the game before. Such Internet game players do not need trust or any normative constraint, such as reciprocity or fairness, to be able to play in orderly ways. Their moves are constrained by the computer program. Their interactions do not have the free form of ordinary life, although freer form can presumably be increasingly programmed into such games. The players form a network that is both gigantic in the number of agents in the network and tiny in the form of what the network manages for them. Some of them might spend forty hours a week in the game, so that it is quantitatively a large part of their lives. It might also take on a large psychological role for them. But it remains highly focused and unlike the melee of ordinary life. It finally bears little resemblance to what we call networks in our lives offline.

At the nearly opposite extreme are self-help groups in which people may bare their personal problems for each other to discuss. Typically, the members of such groups go by names that do not reveal their actual identities. Indeed, in some cases members of such groups have been discovered to be dishonest in pretending to be like the members of the group when they are quite different. While such a group works, however, the members may develop trust relations. They are constantly at risk of being attacked and sometimes even of being exposed, especially if there are Internet savvy members of the group who can ferret out the real world identities of participants.

For most of our dealings on the Internet, trust is not an issue virtually by definition because we face no risk of any significance. In many other dealings, the biggest risk is violation of our privacy, for example, as the result of uncontrollable wider publication of an embarrassing email or, as in the self-help group, the discovery of our actual identity.

5. Enabling Cooperation

In our kind of society, we have cooperative relations within dyads of close friends or relatives; within subcommunities that are small and close enough to be regulated by norms; in exchange relations in varied networks that are regulated by the force of the iterated exchange incentive and its proxy in reputation; and in institutionally protected exchanges. It is only in the first and third of these—dyadic exchange relations and cooperative relationships organized within networks that are typically restricted to specific issues—that trust as encapsulated interest is fundamentally important. For the small community devices of norms and sanc-

tions to work, relationships must be relatively thick. Network exchange relations commonly involve relatively thin relationships that are focused on narrow issues, such as doing reciprocal favors over some range of issues. Dyadic relations can range from the narrowly focused to the nearly all encompassing, as in a good marriage (Franz Kafka might add: if there has ever been a good marriage).⁶

The use of norms grounded in communal sanctions will commonly not work for so decentralized and vast a system as the Internet. A principal reason for this is that the users of the Internet are not also involved with each other in manifold other venues or in overlapping networks, and they do not need to be. I may not even know who is the other person or people with whom I deal; and the other or others may evaporate into the ether at any moment. Potential endgame effects are apt to loom much larger in Internet than in real-world interactions. Endgame effects of instant withdrawal are likely to be relatively rare in small community relations and even in friendship relationships, although there can be sudden blow-ups that lead to complete breaks even in such relationships. I may be involved offline with many of those with whom I deal over the Internet, but not with most of the people who might use the Internet to abuse my interests in some way. For the former, the norms that govern our interactions will often typically precede the Internet in our relationship or will have transcended our relationship over the net.

The Internet is the most extreme version yet established of an exceedingly large anomic society. In its Greek origins, anomic means lawless; in contemporary English it more often means normlessness. Much of our use of the Internet is normless. Moreover, almost all of it is lawless. One can commit fraud over the Internet, and that would be subject to legal action. But many people spend more than a thousand hours a year on the Internet and are never involved in anything that would rise to the level of legal scrutiny—unless they happen to download copyrighted music or videos. In its normlessness and anonymity, the Internet has some of the characteristics of urban life during those periods when one is anonymously walking the streets and meeting no one whom one knows. But it goes far beyond such anomic interactions in that it often actually involves exchanges and even some degree of dependence on one another. Very little, if any, of our standard vocabulary for describing and explaining human relationships fits many of our relationships on the Internet.

Nevertheless, the Internet is enormously valuable to us. It enables us to do many things that would otherwise be extremely difficult or even nearly impossible. In my milieu, for example, I can use the Internet to find out things that it would take days to discover through traditional means, such as library researches. It provides—or is—a novel form of social capital. Often when I am seeking a fact on the Internet, the sources on which I stumble turn out to be the same sources I would have consulted in traditional searches. Sometimes, however, they turn out to be the websites of industrious hobbyists whose catalogs of facts and stories are a labor of love. Once I do find out something, however, I can then very efficiently double-check it in traditional ways or by consulting

⁶ Kafka 1967, entry for 17 October 1921, held the idea of marriage to be beautiful, but he doubted that any instance of it was very good.

traditional sources—such as newspapers and encyclopedias—over the web. For me, the Internet is—apart from its email function and its capacity to deliver documents instantly—primarily a device for discovery of facts, but I am not finally dependent on it for verification of facts.⁷

6. What the Web Enables Us to Be

It would be interesting to draw up a catalog of the peculiarities of the life of the Internet. One is discussed very entertainingly by Sherry Turkle. Turkle says of herself that, on the Internet, she is a “multiple, distributed system”. She is a clinical psychologist, social scientist, author, professor, and cyberspace explorer, where she travels under many identities or names. She could do, as many of those she has studied do, far more and travel under essentially untraceable names that could let her act out any persona she chooses over the anonymous Internet. We cannot do that as readily, if at all, in real life.⁸ Therefore her title, *Who Am We?* She supposes, however, in the contemporary post-modern vision that we are, in some meaningful sense, all of these identities, that we are multiple in nature (Turkle 1996).

Turkle has been characterized as the anthropologist or ethnographer of the Internet, although it might be more accurate to see her as a psychologist of the Internet. She argues that “computers are not just changing our lives but changing ourselves” (Turkle 1996). In a sense, she is generally interested in answering the question what—or who—the Internet enables us to be. She focuses on what the Internet does *to* us. I am here more generally interested in what the Internet enables us to do. I want to understand what it can do *for* us. This focus pushes us to study the Internet as a form of social capital. We can invest in that capital, as millions of people and institutions are doing, thereby increasing the value of the Internet to others, and we can put it to use, as even more millions are doing.

Turkle has interviewed many players of online games, mentioned earlier, to determine what the playing, usually anonymously, does for them when there are well over a million regular participants over a long period of time from many societies. Those facts alone make it sound like an important social phenomenon that we might want to understand. If we are to determine the larger social effects

⁷ If one uses anything from a hobbyist’s website, verification can be a major issue. Websites and online message boards recently have defended Kobe Bryant against sexual assault charges with extraordinary venom directed at the woman making the charge and everyone associated with her. It is exceedingly unlikely that these Internet-assisted libelers have any knowledge on which to base their claims. They have only energy and partisan commitments. Getting caught up in the thrall of a sports fan evidently reduces a person’s IQ by approximately 23 points. Because news networks and television stations want their own websites to have many hits, they put in links to the irresponsible sites. CBS, for example, can get thousands of hits while keeping itself formally out of the libelous debate, which it merely enables others to enter with greater ease (Roberts 2003).

⁸ One of the most remarkable cases of adopting an identity is John Howard Griffin (1960) who had himself blackened so that he could travel the South in the United States in 1959 in order to experience the life of racial discrimination.

and significance of this phenomenon, we need to know who the people are and what are their larger roles in the society. Or in Turkle's language we need to know more about the people that they are outside their personas as players in the game.

Those who are interested in the phenomenon of television in contemporary lives do studies of such things. Television watching similarly occupies a large part of many lives, upwards of forty hours each week. (That is hard to imagine, because it is equivalent to *six hours every day*, say from 6:00 until midnight every evening.) For some of these people television is not a substitute for life in other forms but is merely an entertainment during lives that would otherwise be even more empty and lonely. For many, however, it is thought to be a major alternative to living in supposedly more interesting ways. Robert Putnam thinks that the distractions of television have substantially reduced the social capital that comes from participation in various social groups and organizations (Putnam 1995; 2000).

Eventually, we should want to have similar accounts of the effects of living life on the Internet. Presumably, Internet gaming does some of what television does to social participation. Quite possibly, however, this claim is wrong, because the gamers might be people who are only loosely connected to the kinds of groups and organizations that interest Putnam and other critics of television. Maybe, in Putnam's metaphor, they would be bowling alone with or without the Internet, and many of them may be especially attracted to the Internet in large part because it allows them to bowl alone and anonymously. Perhaps they would not be more socially active even if the Internet collapsed. Of course, it could be the case that computer usage from an early age has created part of the personalities of those who are now turned off of standard categories of social participation. What we need, if we are to understand such worries and loose claims, is hard studies that go beyond anecdotal stories of particular individuals, as interesting and suggestive as these are.

7. Internet Pneumatique

Among the most striking aspects of the Internet is the ease with which one can write letters and thereby maintain relationships across massive distances and long times. I am involved with a fairly large number of email correspondents, but for us the Internet is merely a convenience. It is true that the convenience may be so dramatic that it changes our relationships, much the way the relationships of Parisians in the time of Marcel Proust or of Czechs in Prague in the time of Franz Kafka must have been changed by the fact that they could exchange several letters in a single day by use of the pneumatique. Such quick exchange may tend to heat up the contents of the exchanges and therefore the relationships themselves.

The ease and alacrity of email also changes the scale of our correspondence. I receive and write upwards of a dozen real letters a day in exchanges with real associates, and in addition I get nudged several times daily (I do not count

these as real letters). Before email, I wrote and received far fewer letters. And I very seldom received letters from people I did not know. After being out of the country for more than a week some time ago, I returned to find more than three hundred real messages (after deleting the roughly equal number of junk emails). I did not quickly make it through that backlog and some of it may have died of neglect. There are probably people in the world who are offended at my inattention. Yet I spend more than an hour daily doing email—there were few days before email when I spent so much time on such correspondence.

For Turkle, Internet romance via email is a major topic. For many of us, presumably it is not a major part of our lives, but for some it apparently is. It is possible that the distancing of the Internet is actually an enabler for initiating romantic attachments because one can come to know a lot about another person's sense of humor, intellect, and character more generally in a remarkably safe venue in which it is relatively easy just to drop the correspondence and the correspondent at will.⁹ If the relationship seems appealing, it can always be transferred to the real world. I am told by someone whose authority is probably great that email romancers know that they should attempt a meeting soon—before they develop a strong attachment to someone they would never have liked in person. Without such a meeting, they are at risk of falling in love with an ethereal persona, not with a person, but eventually suffering the anguish of a broken romance with a real person. Email, weirdly, has returned us to the era of the troubadours, with their fancifully embellished beloveds admired from a great distance.

This is a context in which a person might have something at risk so that there is at least a possibility that trust would arise. The risk is that, in order to pursue the relationship further, one must eventually reveal one's identity. Indeed, such revelation is likely to be an element in demonstrating one's trustworthiness and, therefore, in establishing a trust relationship. Again, for trust to be at issue, there has to be something at

risk that will be lost in the event of untrustworthy behavior by the trusted person. For many of us, especially in our lives in business, the professions, and academia, we are more or less fully exposed in many ways on the web. We have websites that include lots of information, some of which could readily be linked to other sources to fill in many gaps. One could say—in the vacuous vernacular—that one trusts everyone else with all of this information. It would be far more meaningful to say that one thinks it worth the risk to put out so much information, because it enables one to make connections that would be valuable. For academics, easing the way for others to read our work is likely to be beneficial to us. There may be heightened risks (for example, of plagiarism or vicious criticism) from such easy exposure, but they seem likely to be outweighed by the benefits.

⁹ Such romantic involvements might be risky in ways that only the Internet can threaten. A couple of years ago, one of the Stanford University servers was reported to have randomly readdressed emails of people with last names beginning with the first three letters of the alphabet. Boastful emails about affairs have been re-sent to thousands of people, have made it into the international press, and have repaid the boasters with grim repercussions, including lost jobs.

Even while we are exposed in such hopefully constructive ways, we can also lead relatively secret lives on the web by the use of other email names through anonymous providers. To do so is somewhat cumbersome for those who are not very heavily invested in browsing and using the web, because it means checking mail on different systems. For those who are a bit klutzy, this sets up opportunities for embarrassment when using the wrong address fully identifies the writer or chat room visitor.

There is an apparently complex and specific aspect of the regime of email that has not been adequately addressed. Within organizations, people often develop paranoid cognitions and engage in dysphoric ruminations about their colleagues. A reason for this is that often they have too little information about what the colleagues actually think, and psychologically they tend to assume that no information is bad information (Kramer 1994; 1998). What seems like ordinary paranoia is often justified, especially in contexts in which there is a clear or potential conflict of interest or a power differential (Cook et al. 2004, chapter 4). Paranoid cognition is perhaps especially prevalent in international relations, especially in contexts in which there is a serious risk of conflict and even hostility. International relations are therefore perhaps thought to be far more hostile than they are in objective fact.

Paranoid cognition is not itself justified by objective grounds for fearing another, but if there is any reason for concern, that concern is apt to be heightened by communications that are very terse, as email communications commonly are. The medium seems to encourage efficient brevity,¹⁰ and there is not enough time for extensive responses to all our emails. There is not much to read between the lines when there are only one or two lines. The medium typically lacks the expansive graciousness of much of traditional letter writing. It also lacks visual clues to what the correspondent might mean, in particular it misses cues that suggest that the correspondent is smiling rather than scowling. Of course, it also lacks the tones of voice that suggest the mood of a person over the telephone. The use of cryptic funny faces as punctuation marks must often help, as though the writer were saying, "Don't take that seriously". Unfortunately, such faces have become too common to be very revealing and there are seldom scowling faces to give us an accurate measure of real variation in mood. Having no baseline, we evidently sometimes invent one.

Insofar as paranoid cognition and dysphoric rumination are in substantial part responses to the terse quality of email and its common lack of tacitly informative cues, the Internet may often actually foster distrust even when there is no objective reason for it and may exacerbate it when there is at least some reason for it. There is nothing additionally at stake in email correspondence beyond what would be at stake in traditional forms of communication, and yet the form may cause us to respond and act as though there is. Part of that form is its instantaneity.

¹⁰ It also encourages a perverse form of inefficient brevity, as when someone sends a quick reply and then follows it almost immediately with several additional notes to cover neglected points.

8. Concluding Remarks

What makes for trust in ordinary contexts is the richness of our iterated interactions. The richness makes the relationships more valuable so that we want to maintain them; and this is our evidence of each other's trustworthiness. Internet relationships typically are very narrowly focused on what would be merely a piece of an ordinary relationship. In this respect, Internet relationships are more nearly like commonplace commercial relationships, in that they are often not repeated but are one-shot. For commercial dealings, however, we have the devices of the market for establishing and correcting reputations that give us and our commercial partners incentives to behave well. It is that broader set of relationships that secures our commercial dealings. If we have a richer relationship with someone with whom we have Internet dealings, that richness most likely comes from our further interactions offline. There may be a few Internet relationships that become moderately rich, as long-standing self-help groups reputedly do. But the bulk of our Internet dealings that are not merely extensions of offline relationships are too focused and pared down for them to rise to the level of grounding trustworthiness and securing trust.

We could conclude that we need devices for securing reliable interactions on the Internet. Or we could conclude contrariwise that the Internet is inherently not very useful for handling the kinds of problems that dyadic trust as encapsulated interest, small-community norms, offline network connections, and third-party enforcers handle in the offline world. So far, it appears that the latter conclusion is correct. The one area in which trust seems to be at issue on the Internet is in ordinary communication via email. But this is an arena in which most of our dealings are with people with whom we have extensive offline dealings. The ease of using email might facilitate the strengthening of our trust relationships with those *with whom we are already connected* and with those who come to us with third-party commendation. It might also facilitate the quick breakdown of a relationship if paranoid cognition comes into play.

The case for and against the importance—or even the relevance or possibility—of trust on the Internet is still open. If we have even a bit of paranoid cognition about science, we might conclude that the case is generally against trust if no one has yet shown that trust matters. At the moment, what we need if a positive case is to be made is explanations of how trust can arise, be strengthened, and be tested on the Internet. That would include an account of just what kinds of risk we face on the Internet when it enables others to work against our interests. For many of us, trust is beside the point in any case, except perhaps in our email exchanges, because the Internet is merely a form of social capital that we can put to use with sometimes extraordinary results—and sometimes great frustration. It is merely the analog of a very large, extremely well organized library whose content can be scanned almost instantly coupled with an even higher speed and internationalized pneumatique.

In this it is like other categories of social capital. It is also like them in that to use it well requires substantial human capital. I have two research assistants who virtually waltz through the Internet with dazzling ease and grace. It is

human capital such as theirs and of the programmers of various search engines, such as Google, that makes the social capital of the Internet so valuable.

Although it may sound perverse to put it this way, the central issue in making the web be more richly useful to us, beyond its being a quick research and communication tool, is how to create networks on the Internet that approach the richness of the networks most of us have in our offline lives. So far, the Internet, the biggest network of all history, may have almost none of the value of usual networks. Those who have hoped the Internet might enhance participatory democracy and civic life more generally might be disappointed in its performance so far. Many of us may be happy enough with its powers of research and communication. It is not only a new form of social capital, it is also a new form of network. It is at once extraordinarily extensive and remarkably shallow in its reach.

Bibliography

- Carnap, R. (1962), *Logical Foundations of Probability*, Chicago (2nd edition)
- Cook, K. S./R. Hardin (2001), Norms of Cooperativeness and Networks of Trust, in: M. Hechter/K.-D. Opp (eds.), *Social Norms*, New York, 327–47
- / — /M. Levi (2004), *Cooperation without Trust*, book manuscript
- Griffin, J. H. ([1960] 2004), *Black Like Me*, San Antonio
- Hardin, R. (1982), *Collective Action*, Baltimore
- (1995), *One for All: The Logic of Group Conflict*, Princeton
- (2002), *Trust and Trustworthiness*, New York
- (2004), *Hume as Contemporary Political Philosopher*, book manuscript
- Hobbes, T. ([1651] 1968), *Leviathan*, London, ed. by C. B. Macpherson
- Hume, D. ([1739-40] 1978), *A Treatise of Human Nature*, Oxford, ed. L. A. Selby-Bigge and P. H. Niddich (2nd edition)
- Johnson, A. W./T. Earle (2000), *The Evolution of Human Societies: From Foraging Group to Agrarian State*, Stanford, 2nd edition
- Kafka, F. (1967), *Tagebücher (Diaries)*, Frankfurt
- Kramer, R. M. (1994), The Sinister Attribution Error: Paranoid Cognition and Collective Distrust in Organizations, in: *Motivation and Emotion* 18, 199–230
- (1998), Paranoid Cognition in Social Systems: Thinking and Acting in the Shadow of Doubt, in: *Personality and Social Psychology Review* 2, 251–275
- Lewis, D. K. (1969), *Convention*, Cambridge/MA
- Mueller, J. (1999), *Capitalism, Democracy, and Ralph's Pretty Good Grocery*, Princeton
- Putnam, R. D. (1995), Tuning In, Tuning Out: The Strange Disappearance of Social Capital in America, in: *Political Science and Politics* 20, 664–683
- (2000), *Bowling Alone: The Collapse and Revival of American Community*, New York
- Roberts, S. (2003), A Feeding Frenzy on the Web Shoves Sanity Right out the Door, in: *New York Times* 8, 27 July, 6
- Turkle, S. (1996), Who Am We?, in: *Wired* 4.01, January, on the web at www.wired.com/wired/archive/4.01/turkle_pr.html