

Social Capital, Trust, and Economic Behavior

by

David C. Rose
Department of Economics
University of Missouri-St. Louis

Please do not quote

10 October 2007

Copyright © 2007 David C. Rose

Social Capital, Trust, and Economic Behavior

Introduction

For any society to enjoy a condition of general prosperity it must effectuate the gains from specialization. This, in turn, requires an environment that supports a wide range of transactions through which the gains from specialization can be realized. The wider the range of viable transactions, the greater will be the gains from specialization and, hence, the more prosperous a society will be at any point in time. Transaction costs are therefore central to the process of economic development because high transaction costs reduce the range of viable transactions. Institutions are therefore an important factor explaining development history because institutions reduce transaction costs.

Transactions inevitably expose parties to opportunistic victimization and it is well known that opportunism, broadly construed, contributes mightily to transaction costs. What is not so well known is that the larger are the groups within which economic activity such as production occurs (e.g., the larger are firms), the more daunting is the problem of opportunism. This is because specialization necessarily increases the localness of knowledge within production units so the decision-maker will often be the only person who knows the entire action set, including the action that maximizes profit. Institutions can be used to reduce opportunities for opportunism and organizational mechanisms can be used for monitoring to drive up the cost of opportunism. Precluding opportunism through strict routines substitutes one problem for another, since they are antithetical to the efficient, entrepreneurial use of local knowledge. Moreover, the more localized knowledge is, the less likely anyone other than the decision maker will know if the decision chose the most profitable action even if the decision maker is perfectly monitored. In short, the more localized is knowledge, the more likely opportunistic decision making will be beyond the reach of institutional and organizational mechanisms.

This is a fundamental problem, for very large firms (that effectuate Smithian gains) that delegate decision making authority throughout the firm's hierarchy to maximize entrepreneurial direction (to effectuate Hayekian gains) are the *sine qua non* of a prosperous free market society. It is through large group economic activity that societies are able to fully enjoy the benefits of scale economies and specialization, but large group

economic activity invites more opportunism which, in turn, undermines the very transactions required for specialization to occur in an entrepreneurial manner.

To enjoy a condition of general prosperity, something is needed to break out of this development catch-22. Opportunism is the fundamental impediment to the development and operation of a market society capable of producing a condition of general prosperity and trust is the solution to that problem. A climate of trust is antithetical to a climate of opportunism because for *A* to say he does not fear that *B* will opportunistically victimize him is tantamount to saying that *A* trusts *B* not to take advantage him – in other words, *A* regards *B* as trustworthy. It is because opportunism is such a daunting obstacle to the operation of a free market economy that being able to trust others is so important.

It is hardly news that trust plays an important role in the development and operation of a market economy. As Kenneth Arrow (1972, p. 357) has famously stated:

“Virtually every commercial transaction has within itself an element of trust, certainly any transaction conducted over a period of time. It can be plausibly argued that much of the economic backwardness in the world can be explained by the lack of mutual confidence.”

Because the risk of opportunistic exploitation is ubiquitous and worsened by large group economic activity that produces localization of knowledge, trust is a *sine qua non* of effective social interaction. Indeed, virtually no one is prepared to argue that trust is not important (there are a few noteworthy exceptions – e.g., Timothy Guinane) but for social scientists, especially economists, trust is much like the weather – we all talk about it but we don’t do anything about it. There has been surprisingly little effort to model trust. Part of the reason for this is that institutions (formal and informal) work through prudence effectuated by external incentives, so to argue that bona fide trust is important is tantamount to arguing that institutions are not.

Frustration over the inability to solve the development problem in much of the world has produced a new respect for institutional economics, and has led to the rise of the New Institutional School of Economics. This has been a very important development. But it is becoming increasingly clear that getting the institutions right is not enough. Increasingly social scientists, even economists, are coming to grips with the possibility that culture holds the key to understanding why there are such large differences in general

prosperity across societies. Since nearly everyone agrees that behaving in an untrustworthy manner is immoral and morality is an important part of culture, it follows that if we think trust is important we probably cannot avoid incorporating culture into our paradigm if we are serious about solving the development problem.

I realize that by saying culture matters I am wading into dangerous waters. This is true for a number of reasons. David Landes has written about the dangers of appealing to cultural explanations, specifically the danger of being charged a chauvinist or racist if one's findings are not politically correct. There are even deeper concerns. First, the economist's version of *deus ex machina* is dropping an explanation into a utility function, and explaining an economic phenomenon in terms of exogenous culture looks a lot like dropping an explanation into a utility function. I am very sympathetic to this objection but just because culture has been used in the past to pull exogenous rabbits of the hat does not mean we must treat culture as unchanging or exogenous in the future. Second, the duping of Margaret Mead shows how easy it is to find what you are looking for when you begin your search with culture. This is partly the result of the third problem, which is the imprecision of the term.

Recently an idea has come along that lends a little more empirical precision to the idea of culture. That idea is the concept of social capital. Few ideas have had as immediate and widespread an impact on the social sciences as the idea of social capital. The purpose of this paper is to review the concept of social capital and explain why it presents a new approach to understanding how people can create an environment of trust in social groups and, in so doing, make viable the largest possible set of transactions through which the gains from specialization can be realized.

First I briefly review the concept of social capital and the concept of trust. I then identify some unresolved issues in the social capital literature, particularly as it relates to trust, and suggest refinements. I explain why the kind of trust required to overcome the local knowledge problem noted above must have certain properties that are not exemplified in existing models of trust from either social capital theorists or economists. As a result we have the right model for trusting behavior in impoverished societies and the wrong model for trust in prosperous societies. Finally, I offer a new definition for social capital and show that social capital is indeed a *bona fide* form of capital.

Social Capital

The literature on social capital is immense and continues to grow rapidly. Most of the work is empirical, and the absence of a consensus definition and theoretical framework complicates matters significantly. This no doubt accounts for the unusually high number of survey/review articles, collected volumes, and so forth on the subject. I limit my attention to conveying the basic idea of social capital and its relationship to trust and economic behavior. For extensive reviews of the social capital literature see Portes (1998), Sobel (2002), Ostrom and Ahn (2003), and Durlauf and Fafchamps (2004).

So what, exactly, is social capital? There is as yet no generally accepted definition, but I believe I can convey a consensus of its meaning thusly: Some societies are more conducive to social behavior, being able to form and sustain social structures that coordinate behavior at higher levels of social aggregation than other societies. Social capital is the name we give that quality. Social capital is important because it affords solutions to problems that either cannot be solved or can only be solved at very high cost through markets, institutions, or government. This view of social capital comports well with the conventional understanding of the term, as is evidenced by the quote below from *The Economist*. Summarizing the modern responses to the question “What gives rise to the wealth of nations?” *The Economist* (02.22.03, p. 74) states:

“For the past decade or so, sociologists have been pushing one more concept, “social capital” – trust or community, in one of its guises – that is now also being taken up by economists. Crudely speaking, the more people trust each other, the better off their society. They might work more efficiently together, for example. In business trust might obviate the need for complicated contracts, and thus save on lawyers’ fees.”¹

The concept of social capital as it is understood today is most closely identified with the work of James Coleman and Robert Putnam.² Coleman (1988, 1990) laid out the basic theoretical framework in an effort to close the gap between sociological approaches that lack an “engine of action,” as rational action provides economics, and economic approaches that ignore social context. Coleman was particularly interested in the role that social capital played as an input into the creation of human capital.

Putnam, Leonardi and Nanetti (1993) demonstrated the power of the idea in the book, *Making Democracy Work*, which found that measures of civic engagement (their measure of social capital) were strongly related to the quality of local governments in Italy. Then, first as an article then later in the form of an even more influential book, *Bowling Alone*, Putnam made a sweeping case for the existence of a disturbing downward trend in social capital in America since the 1950s (Putnam 1995, 2000). Since the publication of *Bowling Alone* the literature on social capital has mushroomed.³

The enormous attention given to social capital may be a reflection of an increasing appreciation among social scientists for the role played by culture in explaining differences in social, political, and economic performance across societies (Weber 1930; Landes 1998; Harrison and Huntington 2000). Dasgupta (2003) has noted that it is “frequently hard to tell apart writings on social capital from those on culture.” Dasgupta (2003) also states: “Talk of beliefs, and we are drawn inevitably to the notion of culture, which is bound up with the idea of social capital.” It seems clear that some cultural characteristics are more conducive to the creation of complex social structures than others (Banfield 1958; Landes 1998; Fukuyama 1995).⁴ But while growing interest in the idea of social capital appears to be closely related to a growing interest in the role played by cultural factors, social capital has the advantage of being a narrower concept than culture and, when viewed as an extension of the concept of human capital, more amenable to rigorous modeling and empirical testing (Coleman 1988, p. S100-101).

Although social capital is more sharply defined than culture, as I alluded above there is no shortage of definitions of social capital (Durlauf 2001).⁵ The definitions do not offer competing views of social capital as much as they emphasize different aspects of the same general idea. The work of Pierre Bourdieu is not nearly as well-known to non-social capital scholars as that of James Coleman or Robert Putnam, but Bourdieu’s (1986) definition of social capital explicitly accounts for one of the most important distinctions in social capital theory. Consider his definition:

“Social capital is an attribute of an individual in a social context. One can acquire social capital through purposeful actions and can transform social capital into conventional economic gains. The ability to do so, however,

depends on the nature of the social obligations, connections, and networks available to you.” [Bourdieu, 1986]⁶

This early definition is surprisingly perceptive, rich, and durable. In a detailed review of social capital literature, Sobel (2002, p. 139) led his article with a definition that closely followed Bourdieu’s and argued “...[Bourdieu’s] definition of social capital fits easily into strategic models of economic behavior.” We will return to the relationship between social capital and strategic behavior later, as this is a very important issue. The key point here is to recognize that Bourdieu’s (1986) definition really contains two equally important but distinct senses of social capital.

The first sense is consistent with the work of Loury (1977) and Glaeser, Laibson and Sacerdote (2002), which is to think of social capital as a characteristic of the individual. In other words, *individual social capital* is the set of characteristics that enable an individual to be more effective in social settings, including but not limited to what is normally referred to as “social skills.”⁷ The better are one’s social skills, specifically in forging and maintaining social ties within and across social networks and organizations, the more social capital one has and the more productive one will be.

In addition to being a form of human capital (as stressed by Glaeser, Laibson and Sacerdote 2002), Loury (1977) and Coleman (1988) stressed that social capital can act as an input to process of creating human capital. Bourdieu (1986) conveys this idea when the first phrase of his second sentence is matched to the last phrase of the last sentence: “One can acquire social capital through purposeful actions...The ability to do so, however, depends on the nature of the social obligations, connections, and networks available to you.” In any case, whether one stresses that social capital is a form of human capital per se or stresses that it acts as an input to the creation of human capital, a clear advantage of the individual social capital approach is that it is largely based on the theory of human capital, which has already proven its usefulness as a framework for theoretical modeling and its amenability to empirical testing.

The second sense is found in the second part of Bourdieu’s definition, “...the nature of social obligations, connections, and networks available to you.” This aspect of social capital does not pertain to the social skills possessed by individuals but, rather, to the nature of the social environment, specifically the existence of social structures within

which an individual's social skills can be exercised. The implication is that it is of little value to have social skills if there are no social networks within which to "network" or civic groups to lead. This second sense of Bourdieu's definition is more consistent with the work of Granovetter (1973, 1985), Coleman (1988), Lin (1988), and Putnam (1993, 1995, 2000), whose approach to social capital emphasizes social connectedness, social ties, and "embeddedness" as the substrate of social networks.⁸ Finally, trust, which is customarily listed with networks as being one of the two major components of social capital, is clearly subsumed in Bourdieu's reference to "the nature of social obligations."

These two senses of social capital are very different in that one is really an extension of human capital theory while the other emphasizes the study of community-level attributes and aggregate social structures. But one thing they have in common is their recognition of the importance of trust. Trust – particularly trustworthiness – is an important component of individual social capital because a reputation for being trustworthy produces considerable benefits to individuals and improves one's ability to socially interact with others. If one has a reputation for being trustworthy, one is more likely to be employed in higher paying professions and is more likely to succeed as an entrepreneur. An important question is where such trustworthiness comes from. Frank (1988) has shown that it could be rooted in tastes and has stressed how emotions assure others that such tastes are real and not the product of strategic mimicry. Finally, the ability to judiciously extend trust is also beneficial, as such individuals are much less likely to become victims of opportunism. This latter ability is emphasized by Toshio Yamagishi who argues that it is a form of social intelligence.

Since Coleman (1988), trust and trustworthiness have been considered key components of aggregate social capital. The basic idea is that behavior within social networks and groups will not be very effective if trust is low. Networks and trust do seem to obviously go together. But the connection between social capital, networks and trust has never been clearly laid out. Recently, however, our understanding of the relationship between social capital, networks, trust, and trustworthiness has been refined. Dasgupta (2003), for example, states: "Social capital, in the sense of interpersonal networks, is certainly necessary if mutually beneficial outcomes are to be identified and the associated agreements reached, but you do not need to know each and every fellow citizen to arrive

at rational beliefs, at a statistical level, about their intended behavior. Trust is the key to cooperation, social capital is merely a means to creating trust.”

Hardin (2002) has argued that trust is simply a rational response to trustworthiness based on encapsulated interests. This approach suggests that social capital in the form of dense social networks exists to create an environment in which repeated interaction makes trustworthiness incentive compatible and therefore makes trusting behavior rational. According to this approach, social capital is not comprised of networks *and/or* trust, social capital is the networks themselves that exist to increase the radius of trusting behavior. The value of networks (and, hence, social capital) is therefore derived from their ability to create trust. There is no doubt that networks have this effect and that this is a valid description of the relationship between social capital, networks, and trusting behavior in most of the world throughout most of human history.

To summarize, while the concept of social capital has improved our understanding of how culture affects social, political, and economic performance, there is as yet nothing that could be called *the* social capital paradigm. An important distinction in the social capital literature is the distinction between individual social capital and aggregate social capital. Trust and trustworthiness are important components of social capital generally and are factors in both the individual and aggregate approaches to social capital. We now take a closer look at trust and trustworthiness.

Trust and Trustworthiness

Many contributions to the trust literature are based on dichotomous distinctions between types of trust (e.g., generalized vs. particularized, thick vs. thin, moralistic vs. strategic, etc.) and in some cases different authors use different terms to convey the same basic idea. In what follows we focus on two distinctions. The first is the distinction between the willingness to extend trust versus the personal characteristic of being trustworthy. The second is the distinction between particularized or personal trust versus generalized trust.

Extending Trust versus Being Trustworthy

The word “trust” often conflates the willingness to extend trust and the characteristic of being trustworthy. In many cases this is not a problem – it is understood from context that the word trust refers to both extending trust and being trustworthy.⁹ But as Hardin (1993,

1996, 2002) has pointed out, the word trust is often used when the idea being discussed is clearly being trustworthy as opposed to extending trust. Moreover, while some theories explain why people are willing to trust others (e.g., Uslaner 2002), other theories treat trusting behavior as simply a rational response to perceived trustworthiness (e.g., Frank 1988; Hardin 2002). Since factors that lead to trusting behavior may differ from factors that lead to trustworthiness, it is important to keep these ideas separate.

So which is the most accurate description of a high trust society: a society filled with people eager to trust others irrespective of their trustworthiness or a society filled with trustworthy people? The former position is taken by Uslaner (2002), who argues that societies that enjoy generalized trust do so because of a belief that extending trust is a moral imperative even in the absence of evidence of trustworthiness.¹⁰ This is sometimes called moralistic trust and is closely related to the idea of altruistic trust. Mansbridge (2001) has argued that altruistic trust, when coupled with an ethic of trustworthiness, can indeed produce a society that is markedly more productive than other societies.¹¹ But what about moralistic/altruistic trust that is not coupled with trustworthiness?

For obvious reasons, the argument that a high trust society can be based on a widespread moral imperative to extend trust has been met with a great deal of skepticism. In addition to Hardin (1993, 1996, 2002), Ostrom (2003) and Levi (1998) are also quite skeptical about the possibility that the extension of trust can lead to trust and trustworthiness. They, like Hardin, view the act of extending trust as morally neutral, a merely rational response to trustworthiness. As such, trust is simply based on the predicted trustworthiness of the trusted individual and is therefore not in any way morally praiseworthy (Hardin 1993). Trust is extended when the trusting party infers that the trusted party is trustworthy because being so is incentive compatible because of “encapsulated interests” (Hardin 2002). According to this view, social capital is interpreted as social structures that exist largely to effectuate encapsulated interests, thereby making trustworthiness incentive compatible and rational and, hence, making the extension of trust rational.

Responding to Hardin’s work, Uslaner (2002) draws a distinction between moralistic and strategic trust. Uslaner clearly views Hardin’s “encapsulated interest” approach to

trust, which is a theory of strategic trust, as the primary rival to his theory of “moralistic trust.” Uslander (2002, p. 15) writes:

“Hardin (2000, 10) argues that claims about the moral foundations of trust are really misplaced claims about trustworthiness rather than trust. But if moralistic trust is based on *presumptions* of trustworthiness, rather than actual evidence, then either Hardin is wrong or the debate is beside the point.”

[italics in original]

A key issue is what is meant by the word “evidence.” If what is meant by “evidence” is evidence that trustworthiness is incentive compatible for the potential trusted party, then trusting behavior is strategic rather than moralistic. But what if by evidence we mean experience has shown that untrustworthy acts are so rare that it would be foolish (indeed irrational) to not presume that others are trustworthy because not doing so will lead to lower material payoffs?¹² In this case, trust would not be strictly relational as is the view of Hardin (2002). With respect to a given individual, it would be based on nothing but hope, but it would nevertheless be rational because it is based on beliefs derived from evidence of trustworthiness for the population as whole. This approach would only be irrational if one did not withhold trust if additional information suggested that the individual under consideration was not trustworthy. This constitutes an intermediate position in that such trust is not moralistic in the sense that one feels compelled to trust out of moral duty, but it is also not strategic in the sense of the term as it might be applied to Hardin’s notion of “encapsulated interest,” yet it is indisputably rational.¹³

There is little doubt that both trust and trustworthiness matter. But it also seems clear that Hardin, Ostrom, and Levi do have a point – trustworthiness is not just more important than trust, it is primary to trust if trust is to be rational. To see this, consider what would happen if trust were present in a given society even though trustworthiness is not. In such a case trusting individuals would constantly be victims of untrustworthy behavior.¹⁴ The effects of operant conditioning alone would likely be sufficient to induce people to stop trusting. It is simply implausible to believe that any individual would continue to extend trust if she were repeatedly victimized by untrustworthy individuals. Even if people did not make such adjustments, those who possessed inherited traits

and/or cultural beliefs that made them the most skeptical about extending trust would soon compete those who were the least skeptical out of existence.

Conversely, if most people were not trusting even though nearly all people in a given society are trustworthy, then the few who were the most willing to trust would enjoy higher payoffs through the benefits of engaging in high trust transactions that those who were unwilling to trust would pass up. Others would observe this and likely figure out that their pessimism is just too costly. Of course, even if people did not make such adjustments, those who possessed inherited traits and/or cultural beliefs that made them the least skeptical about extending trust would soon compete those who were the most skeptical out of existence.

So unless people are irrational, trust without trustworthiness will destroy trust over time. Trustworthiness without trust, however, will likely induce more trusting behavior over time. In short, in the absence of strong institutions to limit opportunism, or when we are dealing with social interaction beyond the reach of institutions, trustworthiness is, at the very least, the *de facto* basis for trust, so the real question becomes this: what explains the existence of trustworthiness?¹⁵ Existing models of rational trust from both the social capital literature and the economics literature model trustworthiness as a rational/strategic response to externally produced incentives. An alternative answer to this question is that it is moral tastes, not institutional and organizational mechanisms that produce incentives that make trustworthiness incentive compatible, that produces bona fide trustworthiness – the kind of trustworthiness that makes the development and efficient operation of a free market society possible because it provides a rational basis for extending trust.

To summarize, trust is viewed here as an amoral, rational, strategic response to the belief that the party to be trusted is trustworthy. Even an immoral person will extend trust to a person who is known to be trustworthy, and only a fool extends trust to someone who is known to be untrustworthy. Extending trust need not be morally praiseworthy for it to occur. Indeed, it seems obvious that if there were a moral imperative to extend trust to those known to be untrustworthy then such a convention (or the society that adopted it) would soon become extinct.

Particularized versus Generalized Trust

Another important distinction in the trust literature is the distinction between trust extended to particular individuals or particular organizations versus trust that is extended to all members of one's society unless there is a specific reason not to do so. Uslaner (2002) characterizes this as the difference between particularized and generalized trust. As Durlauf and Fafchamps (2004, p. 9) point out in the following quote, evidence in the form of expectations of how even strangers will react when extended trust plays a crucial role in generalized trust:

“Sometimes, trust arises from repeated interpersonal interaction. Other times, it arises from general knowledge about the population of agents, the incentives they face, and the upbringing they have received (Platteau (1994a,b)). The former can be called personalized trust and the latter generalized trust. The main difference between the two is that, for each pair of newly matched agents, the former takes time and effort to establish while the latter is instantaneous.”

The idea of generalized trust is more accurately described if one replaces the word “instantaneous” with “presumed” in the quote above. If trust is truly generalized, then the willingness to trust is a foregone conclusion, the consequence of a rule of thumb that says unless there is specific evidence to the contrary, it is rational to presume others to be trustworthy. Fukuyama (1995) and Yamagishi and Yamagishi (1994) have also emphasized the distinction between particularized and generalized trust, the latter referring to this distinction as the difference between personal trust and general trust. Finally, empirical work that looks for evidence of the effects of social capital on society on social, political, and economic performance tends to emphasize generalized trust over particularized trust (Putnam 1995, 2000; Knack and Keefer 1997; La Porta et al. 1997; Uslaner 2002).

Generalized trust is closely related to the idea of thin trust, which is trust that is extended to strangers (Uslaner 2002 equates these terms). But it is important to recognize that thin trust as it is currently defined in the literature is not equivalent to generalized trust. Thin trust is not presumptive in the sense described above; it is still strategic and particular in nature. What makes thin trust “thin” is the fact that it refers to one being willing to trust another with a small matter but not an important one because encapsulated

interests are weak. Such trust is broad (the number of trusting relationships characterized by weak encapsulated interests is large) but shallow (not much will be risked in such relationships). For example, you might trust a stranger with a \$50 bill to make change but not be willing to loan him the use of your car.

Conversely, particularized or personalized trust is related to the idea of thick trust, which is trust that is narrow (it is extended to specific individuals or specific organizations) but it is deep. Such trust is deepest when it is based on mutual affection but it is most often based on incentive compatibility resulting from the expectation of repeated dealings (Bernard Williams 1988). So the closer one is socially to a given individual or organization and/or the greater the prospect for continued association, the more willing one will be to trust an individual.

Particularized trust is strategic by nature. Even trust based on mutual affection is strategic in the sense that one believes trustworthiness to be incentive compatible because it does not improve A's welfare to harm B if A cares greatly about B, so it is correct to say that A can conjecture that B has no incentive to take advantage of A's trust. Repeated dealings, especially as is afforded in dense social networks, is also sometimes viewed as providing a basis for strategic trust in the form of repeated games (Gibbons 2001).

Glaeser, Laibson, Sacerdote's (2002, p. F437) point out that:

“The literature on repeated games (Abreu, 1988; Fudenberg and Maskin, 1986; Kreps et al., 1982) explains why cooperation becomes easier when individuals expect to interact more often in the future. Social connection can substitute for missing, or expensive, legal structures in facilitating investment and other financial transactions (Arrow, 1972).”

In short, both game theory and institutional approaches to trust tend to view trust as a product of social structures. In other words, the role of social capital is to make trusting behavior possible by structuring relationships in such a way as to effectuate incentive compatibility through encapsulated interests.

The distinction between particularized and generalized trust is closely related to another distinction in the trust literature: the distinction between moralistic and strategic trust. As argued above, Hardin (1996, 1999, 2002, 2003), Ostrom and Walker (2003), Ostrom (2003), and Levi (1998) have argued that trust is an amoral, rational, even

strategic response to trustworthiness. Trust has a clear strategic component in the sense that A would not trust B if A believed that B's trustworthiness was not incentive compatible. In other words, A believes that B's trustworthiness is based on the strategic considerations of the possible cost of retaliation (if caught), the cost to reputation, the deadweight cost of a terminated relationship in which there are sunk investments, and so forth. If we think of trust in the context of a repeated trust-honor game such as set forth by Kreps (1990), then both A and B's behavior is strategic.

Trustworthiness itself, however, need not be strategic even if trusting behavior is. According to the encapsulated interest theory of trusting behavior advanced by Hardin, trustworthiness is strategic, but even Hardin acknowledged that trustworthiness might also be normatively based. In other words, A may trust B because A believes that B believes that behaving in an untrustworthy manner is immoral. In this case A need not be motivated by his own morality to base his willingness to trust B on B's morality. Ostrom and Walker (2003, p. 8) appear to concur with Hardin's recognition of this point, stating that: "Hardin argues that to account for trust, one must first account for trustworthiness. Trustworthiness can be seen as primarily based either on incentives or on normative attributes of the decision maker...Trustworthiness is normatively based if the trustee feels morally obligated to fulfill the trust."

The ideas of particularized, thick, and strategic trust all have an important thing in common – they are all inconceivable in large group contexts. This is why they cannot serve as a basis for generalized trustworthiness and, hence, generalized trust. As such, they might all be classified as theories of small group trust. One way to interpret the conventional view that social networks exist to create trust through repeated interactions is that social networks produce the greatest possible radius of trust out of a small group basis for trust (Fukuyama 1995). It seems clear, however, that generalized trust is not merely small group trust on a grand scale (that is, the radius spans the entire society) because the mechanisms through which particularized/strategic trust work (encapsulated interests broadly construed) cannot be extended to the whole of society (which would require extending trust to strangers). If generalized trust exists, there must be something qualitatively different about how it works. This is the root of Uslander's objection to strategic models of trust. Uslander has a point. As noted above, to be truly generalized,

generalized trust must have a presumptive quality to it. But this presumptive quality opens one up to exploitation by untrustworthy individuals. This presumptive quality therefore requires the existence of generalized trustworthiness. Indeed, this presumptive quality is also utterly inconsistent with the rationale behind all small group theories of trust. In societies like the U.S., trust is frequently extended to individuals for whom there is no expectation of repeat play, no opportunity for formal sanction ex post, no reputation to ruin.

Put another way, although Hardin's (2002) notion of encapsulated interest provides a compelling explanation for much of the trusting behavior we observe, it does not provide a convincing basis for generalized trust. If trust is limited to encapsulated interests, then strangers will not presumptively trust each other. Yet in many countries they do, and their willingness to do so correlates with social, political, and economic performance (Knack and Keefer, 1997; Putnam 1993; 1995; 2000). This suggests that while Hardin's approach may explain the most common forms of trust quite well, it still might not explain how the kind of trust that matters most actually works – that being the kind of trust that is required for the development and operation of a market society capable of producing a condition of general prosperity.

Finally, the more unpredictable the world is, the more difficult it is to sustain any form of strategic trust. A world in which there is little change is a world in which nodes of encapsulated interest are most stable. Much of the day-to-day operation of a free market economy is like this, and so encapsulated interests likely provide the basis for a great deal of trust behavior. But in a world of constant change induced by competition with other firms due to their efforts to innovate, adaptation is unending, innovation is unending, change is unending. In such an environment trust based on incentive compatibility through encapsulated interests would be precarious indeed because relationships and circumstances are continuously changing. This is not true for all types of transactions, of course. As noted above, Hardin's approach probably explains most of the trusting behavior that we observe in most of the world and even much of the trust we observe in high trust societies. But it is true of precisely the kinds of transactions that will be shown later to be most closely associated with having entrepreneurial direction of production in large firms. The kind of trust required for enjoying the benefits of

entrepreneurial, not bureaucratic, direction in the context of large firms must rest on something less conditioned on circumstances or the current set of players.

In short, generalized trust is not merely small group trust with a radius of trust that spans the entire society. Generalized trust, it would appear, must be based on something other than incentive compatibility. One possibility is the idea of generalized reciprocity, which Putnam (2000, p. 134) argues is not just the basis for trust, it is the foundation for civilized life. According to Putnam, we all do better when we understand that “The touchstone of social capital is the principle of generalized reciprocity – I’ll do this for you now, without expecting anything immediately in return and perhaps without even knowing you, confident that down the road you or someone else will return the favor.” According to evolutionary psychologists, humans appear to have a natural disposition to reciprocate (Cosmides and Tooby 1992; Dasgupta 2003).

Since reciprocity is a form of trading favors, the idea that generalized trust is rooted in a general disposition for reciprocity is an attractive one to an economist. The idea of generalized reciprocity has merit. But the idea of generalized reciprocity does not overcome the problem of exploitation by untrustworthy individuals. If A trusts B because A feels inclined to repay the favor of having been extended trust by other strangers, then A can be exploited by B if B is untrustworthy. Unless nearly everyone is trustworthy, the repeated exploitation that results will necessarily end the extension of trust to strangers even if one is disposed to doing so out of a desire to reciprocate for trust granted by other strangers in earlier time periods. Something beyond generalized reciprocity is needed to limit such exploitation or else the disposition to reciprocate will be competed out of existence. The only way to solve the problem is for everyone to be trustworthy in the first place, but in that case there is no value-added from the concept of generalized reciprocity.

Generalized trust is more likely a rational response to the existence of generalized trustworthiness per se. Generalized trust can be rational yet non-strategic if the trustworthiness upon which it is based is itself based on moral beliefs that hold that behaving in an untrustworthy way is wrong. Although trust that is based on a moral imperative or a disposition to reciprocate the favor of having been deemed trustworthy by a stranger (generalized reciprocity) is unsustainable because it can be exploited by untrustworthy individuals, a moral imperative of trustworthiness does similarly open one

to exploitation. If I trust you in some way it is easy to see how you might take advantage of that trust to exploit me. But if I never behave in an untrustworthy way, how can you use this to exploit me? One might possibly imagine an elaborate con that takes advantage of one's trustworthiness, but an unwillingness to behave in an untrustworthy way because of having a moral belief that it is wrong is clearly not nearly as easy to exploit as a willingness to trust based on a moral imperative.

Upon closer inspection, it is hard to imagine that generalized trustworthiness could be based on anything other than some kind of moral belief to the effect that it is imperative to be trustworthy. The general problem for any strategic or incentive based approach to trustworthiness is how can A possibly be assured that a *randomly drawn* B will have an incentive to be trustworthy with respect to any transaction A proposes? Any effort to gain such assurance is inconsistent with trust being presumptive in nature. For trust to be truly general, trustworthiness must be truly general, and if trustworthiness is truly general then the issue of whether there is sufficient evidence of incentive compatibility for B's being trustworthy is moot. So if generalized trust is to be rational (and it is hard to imagine that it would last long if it were not), then it must be based on generalized trustworthiness that exists for some reason other than encapsulated interests.

Might trustworthiness itself be based on generalized reciprocity? In other words, might people generally feel that they should be trustworthy to return the favor of having had trust extended to them? The answer is, of course, that they should, but in that case any act of untrustworthy behavior would violate the theory. To the extent that one feels morally compelled to return the favor of having been extended trust with trustworthiness, we are right back to trustworthiness having to be a moral imperative. It seems inescapable that there must be something about the society's culture that makes trustworthiness at least a de facto moral imperative.¹⁶

Economic Approaches to Trust

Economists are naturally inclined to view trust in strategic terms, so for economists the key to understanding trust and trustworthiness is simply a matter of discovering what it takes to make these behaviors incentive compatible.¹⁷ What follows is a review of economic models of trust that, in most cases (but not all; Glaeser, Laibson and Sacerdote 2002, for example), was developed independently of social capital theory.

James (2002) examined the theoretical foundations of the economic study of trust and trustworthiness in the context of a one-sided prisoner's dilemma that is commonly called the trust game, which was formally introduced by Kreps (1990). He argued that there are four basic solutions to the one-sided prisoner's dilemma. The first approach is to solve the problem by changing preferences. Frank (1987, 1988) viewed preferences as a product of an evolutionary process in which some people are randomly assigned an emotional make-up that simultaneously gives them an aversion to cheating (they have a "conscience" so they feel guilty about cheating) and the ability to signal potential transaction partners that they possess such an aversion.¹⁸ Being able to signal an aversion to cheating through emotions (which are defined in such a way as to be involuntary) makes mimicry difficult. In this way emotions confer an economic advantage because potential transaction partners value those who can be inferred to have a conscience. Kandell and Lazear (1992) developed a model of team production in which the team structure is explained as an organizational mechanism that induces trust by increasing how feelings of guilt by strengthening feelings of empathy with one's victims if one behaves opportunistically by shirking duties within the team. Huang and Wu (1994) and Huck (1998) also developed similar models in which the preference for trustworthiness is a function of organizational design.

The second and third approaches considered by James (2002) are explicit contracts and implicit (social) contracts, respectively. The fourth approach is that trust can be based on repeated interaction (see Gibbons (2001, for a review of trust modeled as a repeated game). Models based on contracts of either type (e.g., Ross 1973; Hart and Holmstrom 1987; Kandori 1992; Telser 1980; Spanolo 1999) and models based on repeated interaction (e.g., Axelrod 1984; Klein and Leffler 1981; Fudenberg and Maskin 1986; Abreu 1988; Kreps, Roberts and Wilson 1982; Kreps 1990; Dasgupta 1988; Grief 1993; Neilson 1999) all suffer from the problem of not really capturing what we normally mean by the word "trust." In these models even treacherous people would be trustworthy because it is in their best interest to be so. Robert Frank has referred to this pejoratively as behaving in a merely prudent way. Williamson (1993) is also clearly aware of this problem in pointing out that such "calculative trust" is actually a contradiction of terms. Trust matters most in the absence of any form of assurance from circumstances that

produce incentive compatible – that is, when we have little to assure that it is warranted beyond hope (James 2002).

Zak and Knack (2001) proposed a model of trust based on the fealty of brokers who invest money for investors but possess private information on how well investments actually performed. In their model, investors can invest resources in verifying what brokers tell them or can conserve on such resources by simply trusting them. The structure of their model implies that trust will be a function of formal institutions, informal institutions, and social ties. Social ties affect trustworthiness because reputation costs are posited to be higher the more closely tied a cheater is to his victim. Social ties can be viewed as mathematical distance, which allows their model to explicitly consider degrees of trust. They found that measures of institutional strength that should support trust in their model were clearly related to economic growth.

Zak and Knack's (2001) paper was a big step forward, but the kind of trust they considered in their model is trust that emerges between principals and agents which is a conception of trust that is consistent with Hardin's (2002) notion of encapsulated interests. Zak and Knack's (2001) paper is important in part because it empirically connects a formal model of encapsulated trust to aggregate factors in the form of institutions. There is no doubt that where institutions exist to strengthen the sanctioning of breaches of such trust that this form of trust works better than otherwise. This is a point that has been stressed by Ensminger (2003). There is also no doubt that societies that cannot sustain this form of trust have no chance of being prosperous. The problem with this approach, however, is that there are many societies that have been able to sustain this kind of trust for centuries that are not prosperous (e.g., India). Fealty trust is not what we have in mind when we think of the kind of generalized social trust that is taken for granted in the West but is so rare elsewhere (Uslaner 2002). In short, it is not a model of trust as hope, that is nevertheless rational because it is based on the rational expectation of trustworthiness.

Evidence on Trust and Trustworthiness

In recent years, much work has been devoted to empirically improving our understanding of trust and establishing its importance. Fukuyama (1995) argued that trust is crucial for the development of the kinds of organizations and institutions that are associated with

economic prosperity. Significantly, Fukuyama explicitly argued that generalized social trust is a substitute for the kind of trust that might be found in family firms. La Porta, Lopez-de-Silanes, Shleifer and Vishny (1997) found strong evidence in support of Fukuyama's (1995) thesis in their cross-country analysis of the relationship between trust and the kind of large organizations that are needed to fuel economic growth. Uslaner (2002) documented a precipitous decline in trust. Knack and Keefer (1997) compared measures of trust to economic performance over a sample of 29 countries and found that trust was clearly related to economic performance, although the measures of social capital emphasized in Putnam (1993), which were based on data on membership in social groups, were not related to either measures of trust or improved economic performance. Knack and Keefer's (1997) paper is significant in that its findings strongly suggest it is perhaps a mistake to conflate the ideas of social capital and trust or, at the very least, that we should refine social capital theory to clarify and perhaps emphasize the role played by trust.

There is a disconnect between the empirical literature that examines the relationship between trust and economic performance and microeconomic models that try to explain trust as an exercise in rational behavior. Because these models approach the issue as an exercise in incentive compatibility, trusting behavior is predicated on the specifics of the transaction partner or the transaction relationship. This means that such models are somewhat narrower than they first appear because what they explain is particularized trust based on encapsulated interests and not generalized trust. They are models of small group trust. As a result, the kind of trust considered in existing formal models is a kind of trust that would easily exist and likely be even more important in low trust societies than high trust, prosperous societies. But this has little to do with improving our understanding of the kind of trust that appears to make highly prosperous societies possible; a kind of trust that also appears to be utterly absent in impoverished countries. These models endeavor to explain a kind of trust that has nothing to do with generalized trust, yet it is generalized trust that is actually examined in studies of economic and social aggregates.

This disconnect is further exemplified by the inability of existing models of trust (with the exception of Frank's) to explain the consistent empirical finding from experimental economics that many people are trusting and trustworthy in one-shot

games; games in which there is no organizational mechanism to compel trust, no explicit or implicit contract to compel trust, and no repeat play to make either trust or trustworthiness incentive compatible (Ensminger 2003; Camerer and Thaler 1995; Berg, Dickhaut, and McCabe 1995; Glaeser, Laibson, Scheinkman and Soutter 2000).¹⁹ Indeed, McCabe, Rassenti and Smith (1996) found that almost 50% of subjects attempt to cooperate even in a one-shot prisoner's dilemma game in which players were told they will be playing with a randomized partner only once. The only conclusion that could be afforded by most game theory models of trust in response to this finding would be that about 50% of the population is either irrational, stupid, or both. It seems much more likely that some kind of behavioral rule of thumb was being followed as a matter of practical conservation of cognitive resources, or a social convention was being followed as a matter of moral principle so the question of whether to act in a trustworthy fashion was never really an open one.

More recent evidence from experimental economics appears to support the view that trust may be less strategic than is commonly assumed in game theory models. Cox (2002, 2003) used a triadic experimental design to separate trust and reciprocity from altruism. He found that in the strong social context variation of an investment game experiment, "67% of the amount returned by second movers to first movers can be attributed to other-regarding preferences." (Cox 2003, p. 25) In a recent paper by Carter and Castillo (2003) they devised a set of experiments that allowed the separation of the effects of trust and altruism in a random sample of South African communities. What they found was that while these effects are related, they are nevertheless quite distinct.

In general, we find higher levels of trust across cultures, much higher than would be predicted by Nash equilibrium. Why? I would argue that this is indicative of the ubiquity of small group trust. Economic experiments regarding trust involve small numbers of individuals, usually only two players. Even if this other player is anonymous, this frames thinking in a small group way. As but one way in which small group thinking might be at work, the marginal effects of A's actions in such experiments are not diluted so A empathizes with B even if B is anonymous, so A knows he will feel guilty about harming B if he acts opportunistically. Trust in this context may therefore only tell us that empathy effects are enough to sustain higher levels of trust than predicted by strategic

models. Such findings tell us nothing about the presence or nature of generalized trust. For example, an experiment in which person B interacts anonymously with person A that gives B the opportunity to promote his/her welfare at the expense of A to earn ten more dollars by behaving in an untrustworthy manner will likely result in many more people behaving in a trustworthy way than an experiment in which person B can promote his/her welfare at the expense of *insurance company* A to earn \$100 more dollars by behaving in an untrustworthy way.

The finding in experiments in which other-regarding preferences can be disentangled from incentive effects suggests that at least some forms of trust come quite naturally to us through our desire to promote the welfare of others. This is hardly surprising. We are a small group species so it is only logical that we instinctively look out for each others' interests. Could this be the factor that accounts for generalized trust in the West and Japan? I don't believe so. In fact, I believe this form of trust is quite different from the kind of trust that is needed to support economic development.

In short, I am concerned here with rational trust that is based on trustworthiness that is not the result of specific institutional arrangements or incentive compatibility resulting from repeat play or mutual affection but, instead, that is the result of a belief that the individual being extended trust believes that being trustworthy is a moral imperative. This conception of trust is consistent with what James (2002) refers to as "trust as hope" as opposed to "trust as prudence."²⁰

To summarize, generalized trust is necessarily based on a norm of trustworthiness. But this cannot be based on generalized reciprocity or any small group theory of trust. It is simply not possible for a norm of trustworthiness to arise from strategic models of trust. The evidence from economic experiments suggests that while strategic factors may be sufficient, they are not necessary for either trust or trustworthiness. At least some individuals in every society have either hardwired or learned moral tastes to support trustworthiness. In some societies it appears that nearly every individual is trustworthy. Since hardwired moral intuitions do not vary from society to society, this can only be the result of learned beliefs or moral tastes that produce a conviction that always behaving in a trustworthy manner is a moral imperative.

Human Capital and Social Capital

Taking Capital Seriously

Although social capital is a tremendously important idea, there are a number of problems with the concept. As noted above, there is no shortage of definitions of social capital. A common feature among them is their functionalism. Most definitions fail to distinguish what social capital is from what it does.²¹ Sobel (2002, p. 146) notes that both Coleman and Putnam are guilty of this practice, citing Coleman (1988, p. S98) thusly: “social capital is defined by its function.” The definition offered above as a consensus is similarly functional.

The prevalence of functional definitions in social capital theory has created a problem for much of the empirical literature by producing circular arguments of the following sort: a society is conjectured to be successful because it is observed to have social capital but the social capital itself is taken to be evidence of such success (see Sobel 2002, p. 146, who cites Portes (1998) and Durlauf (1999) as having made essentially this point). The result has been an enormous empirical literature (see Durlauf and Fafchamps 2004 for an excellent review) that equates social capital with the outcomes it is purported to produce. As a result, researchers are always successful in finding that social capital produces social benefits because they look for positive social outcomes (benefits) and then call that social capital or evidence of social capital.²² Indeed, the importance, even existence, of social capital is never treated as an open question.

Even if this circularity were avoided by the adoption of a non-functional definition of social capital, there would still be the problem of empirical imprecision. Solow (2000) has stressed the problems with measuring social capital, questioning whether it is even properly viewed as a form of capital. Dasgupta (2003) has referred to it as “fiendishly difficult to measure.” Manki (2000) has gone so far as to suggest that social scientists avoid using the term. Bowles and Gintis (2002) have stated that to avoid confusion, the expression “social capital” should be replaced by the concept of community, which “focuses on what groups do rather than what people own.”

Arrow (2000) has argued that social capital doesn’t possess characteristics that we associate with physical capital. According to Arrow, physical capital has three basic characteristics: extension in time, sacrifice in current periods to enjoy benefits in future

periods, and alienability. While social capital may possess the first characteristic, Arrow argued that it does not possess the last two. Solow (2000) also argued that the use of the word capital is suspect. Both Arrow and Solow's arguments appear to refer specifically to aggregate social capital, not individual social capital. Their criticism is fundamental but I submit that their conclusion may be the result of adopting an implicit definition of capital that is, perhaps, unnecessarily narrow.

The criticisms leveled by Arrow and Solow should concern any economist, particularly those who believe the idea of social capital to be a fruitful one. But arguing that X is not Y because X does not share many of the characteristics of Y is not the same thing as showing that X and Y do not both belong to the same set. Similarly, arguing that social capital does not appear to have characteristics that are evident in physical capital is not the same thing as showing that social capital is not capital. The adjective "physical" in front of "capital" implies that physical capital is a proper subset of "capital." Similarly, the adjective "social" in front of "capital" implies that social capital is a proper subset of "capital." Obviously, the set "capital" can have two proper subsets that are completely disjoint. No one is arguing that social capital is physical capital.

I adopt a more fundamental definition of capital here than that implied by Arrow and Solow's criticisms of social capital. In my view, capital per se is that which is used in production of something useful that has the property that it is nevertheless not "used up" in the process. When building a house, a hammer is capital but nails are not, because the hammer is still there when you are finished but the nails are not. When building a house the saw is capital but labor is not, because the saw is still there but the labor is gone – indeed perfectly sunk. When running a firm, knowledge of how the production process works is capital (human capital) but time spent solving problems is not, because that knowledge does not disappear after using it.

This approach to defining capital is rich enough to capture the essence of what differentiates capital from labor and intermediate goods or inputs. Consider baking cakes. There are three categories of things that go into baking cakes. The first is the oven, baking pans, bowls, mixer, etc., which any economist would recognize as capital. The second is the eggs, flour, sugar, etc., which are ingredients and which any economist

would recognize as throughputs. Finally there is labor, which uses capital to transform throughputs (the ingredients) into final output – cakes.²³

These three categories are in some sense irreducible. There is a proportional relationship between throughputs and output. Indeed, the ingredients are “in” the cake, as it were, although they have been transformed by labor with the help of capital. Labor is necessary in that it effectuates this transformation, but its relationship need not be proportional and generally is not (at least not in pin factories). Like throughputs, labor is used up, but unlike throughputs it is nowhere to be found in the final cake (what we see is evidence of labor by virtue of the transformation of throughputs, but that is not labor itself). Capital either makes possible or lowers the effort required by labor to transform throughputs into the final output. But unlike throughputs or labor, capital is still there to be used again after the cakes are made. The only sense in which capital is used up is that it is worn out (depreciation), but this fundamentally different from the sense in which throughputs are used up (they are in the cake) or labor is used up (it is gone forever – a perfectly sunk cost). One might nevertheless quibble that depreciation is just another way of being used up, but this, we will find out shortly, only proves to strengthen the arguments that follow.

To reiterate, capital is anything that is used but not used up in the production of something. This definition is very simple, very fundamental, and it easily accommodates both physical and human capital as proper subsets. So is social capital a valid form of capital by this definition? Put another way, is social capital also a proper subset of capital? To answer this question we will examine aggregate social capital and individual social capital separately. We begin with aggregate social capital, specifically the question of whether social networks are indeed a form of capital as defined above.

Social networks are as central to social capital theory as trust is. Granovetter (1974, 1985) has stressed the importance of weak social ties that are the substrate for social networks. Social networks are also central to both Coleman’s and Putnam’s approach to social capital. Dasgupta (1999, 2002) has argued the most strongly for the importance of social networks to social capital theory in an effort to avoid the functionalism of existing definitions of social capital. Specifically, he states: “...social capital is most usefully viewed as a system of interpersonal networks.” (Dasgupta (2002, pp. 6-7)). Durlauf and

Fafchamps (2004, p. 26) put it nicely: “Dasgupta argues that social capital should not be defined in terms of the presence of cooperation or some other outcome; rather that it should be regarded directly as social structure.” Durlauf and Fafchamps (2004, p. 55) continue: “[Brehm and Rahn’s (1997)]...finding is indicative of the empirical importance of Dasgupta’s (2002) argument that social capital should be modeled as a network.”

Networks resulting from social ties that were cultivated by prior investments in personal and professional relationships certainly look like a form of capital.²⁴ Networks appear to satisfy the first and second characteristics that Arrow identifies: extension in time and sacrifice in present periods to enjoy future benefits. The only possible problem is the third characteristic of alienability. In my view, the fact that physical capital is by nature alienable does not mean that alienability is an essential attribute of capital generally. Most basketball players are tall but that does not mean that tallness is a necessary characteristic of an individual who defines himself as a basketball player. Indeed, not all basketball players are tall. Suppose that for some reason ovens suddenly became inalienable. Would ovens then cease to be capital if they could still be used to bake as before? If ovens were no longer capital after becoming inalienable, then what, exactly, would they be? The fact that alienability is a feature of physical capital in economic models does not compel us to regard alienability as a necessary feature of capital in general. Finally, social capital as it exists in a social network would seem to be no less alienable than human capital generally, but I doubt that Arrow would claim that human capital is not a form of capital because human capital is inalienable.²⁵ To be fair to Arrow, his comments were directly specifically toward physical capital. But then what this means is that he has demonstrated that social capital is not physical capital. We are only concerned with whether social capital is not a form of capital, not whether it is a form of physical capital, which it obviously is not.

Knowledge stored in the mind of an individual of how to perform a task in a production process is capital – human capital – and by my approach to defining capital it is actually a purer form of capital than a hammer, a machine, or any other kind of physical capital. It is purer because not only is it not used up as it is used, it doesn’t even depreciate with use; if anything, it appreciates with use!²⁶ Social capital in the form of a social network is ultimately stored in the minds of the individuals who comprise the

network. The relationships that comprise a network are ultimately sustained as memories of the experiences of past dealings with individuals and this set of memories about others (and they of you) is what actually lowers the cost of coordination. These memories are a physical reality – they are stored in a physical brain – there is no mind-body Descartes dichotomy at work here. Such information is generated by experience with others and initial experience is not likely to have been very profitable because initially trust may be low.²⁷ Over time, however, the value of relationships becomes clearer. In this sense we invest in relationships with others in our minds through accumulated knowledge in the form of memories and they do so as well in theirs. Like human capital, the exercise of this form of capital strengthens it, literally strengthening the neural connections that connect and sustain memories; it does not in any way wear such memories out. In this way even aggregate social capital in the form of social networks is itself a product of human capital and, like human capital generally, it is a purer form of capital than physical capital.

We now turn to individual social capital. Is this properly viewed as a form of capital? Existing treatments of individual social capital (Loury 1977; Glaeser, Laibson and Sacerdote 2002) certainly view individual social capital, in the form of social skills, as being a form of human capital. If one accepts that individual social capital is human capital, then it is necessarily a form of capital if we view human capital as a legitimate form of capital. But whereas the potential problem with the phrase “aggregate social capital” was the word “capital,” with the phrase “individual social capital” the potential problem is the word “social.” In other words, that individual social capital, if it exists, is capital seems obvious enough because it is clearly a form of human capital. The real question is whether there is enough of a difference between the concepts of human capital and individual social capital to warrant a distinction. The answer, I believe, is yes, and lies in understanding the role played by the personal characteristic of trustworthiness in the relationship between human capital, individual social capital, and aggregate social capital. We will return to this issue in the next two subsections.

Finally, regarding Solow’s criticism about the immeasurability of social capital, this is certainly a valid criticism of existing empirical work. Perhaps social capital is an idea that is like many important ideas in science – its real value is that it gets us thinking in a

fruitful direction. There is little doubt that much of our beneficial social behavior is dependent on the existence of phenomena that can only be understood if investigated at a social rather than individual level of analysis.²⁸ Perhaps social capital is better viewed as a paradigmatic concept that plays a role analogous to that played by the concept of utility in neoclassical economics. No one thinks he/she can measure utility (and least not yet) and most economists do not believe it is necessary that a cardinal measure of utility must exist in principle, yet utility is the core idea upon which the neoclassical model is based.²⁹ Why can't social capital play an analogous role in the social sciences generally, generating secondary and tertiary theoretical implications that, in turn, generate empirical implications that *are* testable? Whether social capital is a fruitful concept for social scientists can thus be answered in the same way that this question was answered for utility: if empirical evidence repeatedly suggests that models based on the social capital framework are useful, then we will gradually come to accept that social capital is indeed a useful construction.

Avoiding the functionalism of the consensus definition I gave above, I now define social capital thusly: *The relationships, networks, and individual characteristics that improve the effectiveness of social interaction and that are the result of prior investment.* This definition is based on what social capital is, not what it does. It captures the ontological properties of the “extension in time” component and the “sacrifice in current periods to enjoy benefits in future periods” component that Arrow attaches to physical capital. It does not capture Arrow's alienability component, but neither does human capital so we can assume that alienability is an attribute of physical capital, not capital in general. In a free market society you cannot sell your human capital even if you want to because others are not allowed to purchase it because slavery is forbidden. Others can purchase a flow of labor services from you, of course, but that is not purchasing you, it is purchasing a flow of services through time, which is more accurately interpreted as renting human capital than owning it.

Trust as an Input to Social Networks

As noted earlier, increasingly the social capital literature treats trust as a rational response to trustworthiness and trustworthiness as an output of social networks. In short, the repeated interactions afforded by dense social networks make trustworthiness incentive compatible and, hence, makes the extension of trust in such “circles of trust” rational. There is no doubt that social networks have this effect at the margin. There is abundant evidence that high levels of trust lead to superior social and economic performance. But if the key to creating and sustaining trust is networks, then why is it apparently so hard to create and sustain such networks? In what follows I address this question and, in so doing, connect individual social capital to the idea of aggregate social capital by treating trustworthiness as an input, not an output, of social networks.

As currently modeled and discussed, trust generated by networks is, by nature, small group trust in that it is based on encapsulated interest. Networks basically create more trusting behavior than would otherwise exist by extending the radius of small group trust. But what about large group trust, particularly generalized trust? Does the existence of generalized trust change the nature of the relationship between social capital, networks and trust? Might trust, particularly generalized trust, be better viewed as an input to social networks? An alternative approach that I propose here is to view networks not as deriving their value indirectly as inputs to creating trust that makes cooperation possible but as deriving their value directly from facilitating information sharing between individuals and between organizations.

It is well known that information sharing even between competing firms is common and can be socially beneficial (Powell 1990). Information sharing is a way that we, as a society, derive maximum benefit from balancing cooperation and competition.³⁰ Information sharing also benefits the individuals and organizations involved directly by lowering the risk of *relative* failure substantially. What looks like collusion against idiosyncratic failure can also have the effect of increasing innovative activity in equilibrium by reducing its downside risk. Powell (1990) and Powell and Brantley (1992) have shown that there is extensive sharing of information in the biotechnology industry even between competing firms, and this practice is encouraged by firm owners.

But sharing information is risky. It therefore requires extending trust, and such information would not be shared with parties known to be untrustworthy, so cooperation in the form of information sharing is a consequence of perceived trustworthiness. That is, trustworthiness makes an information sharing network possible, it is an input to the network, as it were. This interpretation is consistent with any network because the primary function of any social network is the sharing of information. Trustworthiness is not a consequence of networks, networks are a consequence of trustworthiness.

Because of economies of scale the larger a network is the more effective it is, *ceteris paribus*. Large networks are therefore superior to small networks. But the larger is any social network, the less likely that small group trust can sustain the network. Very large networks and very large organizations are therefore possible only if generalized trust exists, so generalized trust is an input, not an output, for large social networks and organizations.

Perhaps the current approach of viewing social networks as mechanisms to extend the reach of particularized trust derived from encapsulated interests gets the story basically right about most social networks in most of the world, but gets the story wrong in the kind of high trust societies that are capable of producing a condition of general prosperity. In prosperous societies, perhaps networks are not created to sustain trust based on incentive compatibility arising from repeated dealings or institutionalized sanctions but, instead, are created to effect social gains directly (e.g., sharing technological information in a loose consortium or sharing information about employees in job markets). This raises the following question: which kind of trust is the more important input for large networks and organizations: small group trust or large group (generalized) trust? Generalized trust would obviously make it possible to construct extensive networks within which repeat dealings are infrequent (e.g., a network of professors who write letters of recommendation for students). Again, this argument is in no way inconsistent with the current view. Indeed, it seems likely that the Dasgupta-Hardin mechanism is always operative (all societies have small group trust) while the mechanism I describe is only rarely operative. What makes this latter mechanism important is that while it is rare, it may be essential for fostering economic growth and development if growth and development require large networks and organizations (which we will address later). In

short, rapid economic growth and development might be rare in part because large networks and large organizations are rare, and these are rare because generalized trust is rare. In any event, it is ultimately an empirical question whether trust is more of an input to social networks in high trust, highly prosperous societies and more of an output of social networks in low trust, poor societies.

In the next subsection I focus on the characteristic of trustworthiness to show that it is a form of human capital, specifically a form of individual social capital, and illustrate how it connects individual social capital to aggregate social capital. In so doing, I establish that this aspect of individual social capital is sufficiently social in nature to warrant being differentiated from human capital.³¹

Reconciling Individual and Aggregate Social Capital

Trustworthiness is a personal characteristic that is a part of an individual's stock of human capital because it benefits an individual to have a trustworthy reputation. Like any form of human capital, trustworthiness is not used up as it is exercised. When A trusts B and B behaves in a trustworthy manner, this in no way reduces B's ability to behave in a trustworthy manner with respect to C either now or in the future. Moreover, the greater the number of transaction partners who can claim that B has never treated them in an untrustworthy manner, the more willing individuals will be to trust B.³²

Since these benefits are realized through social interaction, generating "returns from interactions with others," trustworthiness can also be viewed as a form of individual social capital (Glaeser, Laibson and Sacerdote 2002). The transaction partners of a trustworthy individual also benefit from not being cheated. Trustworthiness therefore generates a positive externality in the form of a spillover to the transaction partners of the trustworthy individual. According to Glaeser, Laibson and Sacerdote's (2002) approach, this alone is sufficient to establish that trustworthiness is not merely human capital or even merely individual social capital; it is in fact aggregate social capital in that it "...incorporates all of the cross-person externalities generated by the different types of individual social capital." But the existence of positive externalities that accrue to transaction partners does not connect human capital to the social structures that we associate with aggregate social capital.

Can we connect individual social capital, via human capital through the characteristic of trustworthiness, to aggregate social capital so they can be reconciled as two parts of a coherent theory of social capital rooted in human capital theory? Ostrom and Ahn (2003, p. xxv) noted that “Building and utilizing physical, human, and social capital are analytically separable, but closely intertwined in reality.” Since individual social capital is already acknowledged as being rooted in human capital, but it is aggregate social capital that makes the idea of social capital so distinctive and that has stimulated so much interest in new ways of thinking about social behavior, it follows that this is an interesting question indeed. As Glaeser, Laibson and Sacerdote (2002, p. F439) point out, this is a difficult question to answer. They write:

In theory, aggregate social capital incorporates all of the cross-person externalities generated by the different types of individual social capital. Hence, aggregate social capital measures social characteristics that yield market and non-market returns to a society. Our definition of aggregate social capital is thus quite close to the usual definitions of social capital. Unfortunately, the path from individual to aggregate social capital is difficult, because of the extraordinary importance of social capital externalities. *The complexity of aggregation means that the determinants of social at the individual level may not always determine social capital at the society-level.* [emphasis added]

We have already established that the personal characteristic of trustworthiness is a form of human capital and, since its benefits are derived from social interaction, is also a form of individual social capital. Now imagine that trustworthiness is unconditional in the sense that individuals do not even consider whether being trustworthy in a particular situation is incentive compatible – they automatically act in a trustworthy manner as a rule of thumb. Suppose that such unconditional trustworthiness is so widespread that it reaches a kind of critical mass, qualitatively changing the social environment, making it safe, even prudent, to presume that even strangers are trustworthy unless there is clear evidence to the contrary. In such an environment generalized trust would naturally emerge because it would be a rational response to the norm of unconditional trustworthiness. The emergence of generalized trust would, in turn, lead to the

spontaneous emergence of social networks and organizations that require a high trust environment. They would indeed emerge at scales that would be unfathomable if trust were based on encapsulated interests or on incentive compatibility achieved through formal and informal institutional sanctions.³³

In short, social capital is capital for a number of reasons. Aggregate social capital in the form of social networks can be viewed as an indirect form of human capital (and therefore capital generally) in the sense that any social network really exists in the minds of those who comprise the network in the form of memories which are, ultimately, physical realities. Beehives exist first in the minds of bees, not the other way around. But the construction of some social networks – those that are key to the development and operation of a market economy – requires trust and trust requires the existence of a norm of trustworthiness. Trustworthiness, in turn, is a personal characteristic that is an element of one's individual social capital and, hence, is an element of one's human capital. So aggregate social capital is capital that is built with another form of social capital – individual social capital – in the form of trustworthiness. Trustworthiness is therefore what connects individual social capital to aggregate social capital by acting as an input to the creation and maintenance of social networks. If this trustworthiness is so widespread that it can be presumed, then generalized trust will exist, making possible social networks and organizations that are so large that repeated dealings would be insufficient to provide incentive compatibility through encapsulated interests. Such networks therefore do not exist to produce trust. Once such a network or organization exists, however, there is no reason why clusters of repeated dealings cannot emerge and therefore small group trust – specifically thick trust – might emerge as well. These two senses of trust need not be mutually exclusive.

To sharpen these ideas further, think of human capital – specifically individual social capital in the form of the personal characteristic of being unconditionally trustworthy – as machine A, and think of a particular social network as machine B. A can produce direct benefits by making x (a good or service made possible by possessing the characteristic of trustworthiness). If, however, there is a critical mass of A machines, then the result will be the spontaneous emergence of B, which is also capital. Just because machine A is used in the production of machine B doesn't make B any less a form of

capital. Machine B increases productivity of many productive activities, so dramatically so in some cases as to constitute an increase in total factor productivity in the sense of Dasgupta (2002). Finally, the existence of machine B might result in an increase in the production of A. The feedback from B to A is analogous to Loury (1977) and Coleman's (1988) view that one important role of social capital is increasing the production of human capital.

To summarize, human capital in the form of the personal characteristic of unconditional trustworthiness makes an entirely new set of social structures feasible that yield benefits at higher levels of social aggregation. When unconditional trustworthiness is so widespread as to be a social norm, it goes beyond being merely human capital that functions in a social context (individual social capital); it becomes properly viewed as a form of *aggregate* social capital in that it makes social structures possible that are themselves part of a society's stock of social capital.³⁴

Summary

Most social capital scholars consider trust to be an important component of social capital. Increasingly, however, scholars have come to view social capital as social networks and not trust per se. According to this view, social networks derive their value by creating an environment in which trustworthiness is incentive compatible and therefore extending trust is rational. I have argued that trust might also be viewed as an input to social networks. Although in most of the world throughout most of human history social networks existed in large part to extend circles of trust, the social networks that exist in high trust societies do not produce trust but, instead, are made possible by a climate of trust. In other words, in high trust societies trustworthiness is the input and the social network is the output. This approach is not inconsistent with the conventional approach since once a network is in place (for whatever reason) it inevitably creates encapsulated interests through repeated dealings and therefore provides incentives for trustworthiness.

Most existing theories of trust are small group theories of trust. The larger the scale over which social structures operate (the more extensive the network or the larger the organization), however, the less frequent will be repeated dealings between the same individuals and, hence, the more important is generalized trust. In other words, the larger

the scale over which we consider trust, the more fluid is membership in the social network, and the more uncertain is the nature of transactions that occur within such networks, the harder it is to imagine that they exist to provide a basis for strategic or calculative in the sense of Coleman (1990) or Hardin (1993, 2002). It is more plausible that large social networks and large organizations have other reasons for existence; that their existence is based on generalized trust.

There are two forms of social capital: individual social capital and aggregate social capital. To the extent that networks are based on trust, aggregate social capital is indirectly based on individual social capital in the form of the personal characteristic of unconditional trustworthiness. In this way trustworthiness connects human capital to individual social capital to aggregate social capital. When a sufficiently large number of members of society are unconditionally trustworthy, generalized trust emerges spontaneously, making large social networks and large organizations less costly to create and sustain. But for this to occur, a norm of unconditional trustworthiness must exist.

The kind of trust that flows from our hard-wired moral intuitions, from institutions, or from incentive compatibility resulting from the expectation of repeat play in social networks provides inadequate support for a modern market economy. A sense of trustworthiness is required that is based on something more fundamental, something capable of giving rise to a climate of genuine trust. Moral tastes derived from moral beliefs that hold that untrustworthiness is wrong seem a likely possibility. In this case individuals extend trust because they believe others to be trustworthy. Others are believed to be trustworthy because they have moral tastes that deem untrustworthy behavior to be wrong and therefore expect to experience high guilt costs if they behave in an untrustworthy manner. This avoids the logical catch-22 associated with assurance derived from prudential restraint. It also provides a basis for a climate of genuine trust that is nevertheless rational because moral tastes are antecedent to strategic decision making. This, in turn, produces an environment with the lowest possible transaction costs and the most secure property rights even in the context of democratic decision making, without sacrificing the assumption of rational decision making.

Endnotes

1. I thank Joe Anemone for this reference.
2. Manski (2000(3) JEL p. 122) gives a brief review of what others have said regarding the origins of the term. Most view it as originating with Coleman (1988) or Putnam (1993). Durlauf (1999) credits Loury (1977) while Glaeser, Laibson, Scheinkman, Soutter (2000) credit Jane Jacobs (1961). Ostrom and Ahn (2003) say the idea can be traced back to de Tocqueville ([1840] 1945) and Hanifan (1920).
3. For a thorough catalog of scale and scope of studies, see Durlauf and Fafchamps (2004).
4. Curiously, Banfield (1958) and Fukuyama (1995) stress the importance of the importance of family in culture as a roadblock to development. The basic idea is that where family is strong, there is less demand for more formal social institutions. The causation might also be reversed: strong social institutions make one less dependent on family and therefore families are weaker in societies that possess strong social institutions – strong social capital.
5. See Durlauf and Fafchamps (2004) for a review of a number of the most influential definitions.
6. Coleman (1988) was apparently unaware of Bourdieu (1986) (Coleman (1988, p. S95) states flatly that “In this paper, the concept of social capital is introduced....” and he does not cite Bourdieu. It is surprising how close their conceptions of the idea come to one another and may be taken as an indication of the validity and importance of the idea, that it was a “multiple” and therefore almost an inevitable next step to be taken in the social sciences given the questions that remained unanswered and the state of the existing paradigm.
7. Individual social capital is not be confused with individual-level studies of social capital. In the former, “individual” refers to a given individual’s characteristics – friendliness, investment in personal relationships, the size of one’s Rolodex (Sacerdote). In the latter, “individual” refers to the level of empirical analysis to contrast with aggregate level analyses.

8. Although Coleman (1988) emphasizes role played by social capital in the creation of human capital, his definition of social capital itself does not comport with the individual social capital approach of Loury (1977) and Glaeser, Laibson and Sacerdote (2002). In other words, he does not view social capital as a form of human capital but, rather, as the connectedness and networks emphasized in aggregate social cost theory. Specifically, Coleman (1988, p. S98) states: “Unlike other forms of capital, social capital inheres in the structure of relations between actors and among actors. It is not lodged either in the actors themselves or in physical implements of production.”
9. If one assumes that the extension of trust is amoral, that it will only occur if it is rational, then using the word trust to characterize an environment in which there is both a high level of trust and trustworthiness is sensible, because the former would not be present without the latter.
10. It should be noted that Uslaner does not argue that moralistic trust requires trust to be extended to everyone, including those who are known to be untrustworthy. According to Uslaner, moralistic trust only requires that unless there is a good reason to believe the individual in question is untrustworthy or some other factor (e.g., extreme temptation that might result from a sack of a million dollars in unmarked bills), it is right and proper to extend trust.
11. Joan Mansbridge (1999, p. 290) defines altruistic trust as trust in which “...one trusts the other more than is warranted by the available evidence, as a gift, for the good of both the other and the community.” Coleman (1990: 304) also notes “...a group whose members manifest trustworthiness and place extensive trust in one another will be able to accomplish much more than a comparable group lacking that trustworthiness and trust.”
12. Although he was speaking about behavior within social networks, this point nevertheless comports well with Dasgupta (2003), who states: “Social capital...is certainly necessary if mutually beneficial outcomes are to be identified and the associated agreements reached, but you do not need to know each and every fellow

citizen to arrive at rational beliefs, at a statistical level, about their intended behavior.”

13. Hardin (2002) views trust as a three part relation: A trusts B to do x. Although trust that is extended on the basis of general statistical evidence is not relational in the sense that A does not have to be assured that B has a vested interest in being trustworthy with respect to A in commission of x, in the framework I developed later, I consider the possibility that in the absence of negative information about B, A's trust of B is not conditional on B in the sense that B has an incentive to be trustworthy, but A's trust of any B is conditional on the nature of x.
14. It might seem obvious that if people were generally untrustworthy they would know better than to trust. But it is possible that most are untrustworthy while believing that most others are not, so extending trust, even by an untrustworthy individual, is still rational. Moreover, it takes only a relatively small proportion of the population to be untrustworthy for those who extend trust to be constantly punished for doing so. Consider leaving a bicycle unattended. If 80% of all people eagerly extend trust (and are trustworthy themselves) so 20% are not trustworthy and therefore don't extend trust (they engage in psychological projection), it only takes a short time in busy area for it to be statistically inevitable that the bicycle is stolen.
15. Note that societies that possess strong institutions to limit opportunism have them for a reason, so the kind of trusting behavior we observe in such societies is suspect if our goal is to understand how trust and trustworthiness supports voluntary transactions in environments in which formal institutions are not effective. Survey data that show that Japanese people are far less trusting even though their behavior appears to be highly trustworthy supports this claim (Yamagishi).
16. Ensminger (2003, p. 189) notes that unlike Coleman (1990) and Hardin's (1993, 505) “calculative” approach to trust, “...to Toshio Yamagishi and Midori Yamagishi (1994, 136), this is merely one form of trust, what they refer to as “knowledge-based trust.”...But Yamagishi and Yamagishi identify a second form of trust that they label “general trust.” This is a cognitive bias based on a belief in the goodwill and benign intent of the partner.” For such a belief to be rational it must not be systematically

mistaken – there must be a rational basis for presuming that others are generally trustworthy and that basis must be that they are, in fact, generally trustworthy.

17. My colleague, Larry White, has suggested that we call this “sub-game perfect trust.”

18. The commitment problem was first elaborated by Schelling (1960).

19. There is an abundance of experimental economics literature that shows that people are quite willing to trust others, even strangers, in one-shot games (Cramerer and Thaler 1995). An obvious question is whether the willingness to trust varies across groups and societies. Ensminger (2003) finds that there are substantial variations and this points to the need for theories of trust that are capable of explaining this cross-sectional variation. But even in the lowest trust societies that Ensminger encountered, there was still far too much trust to be consistent with the predictions of the “trust as prudence” model of trust. One could object to this claim by saying that an implicit social contract to be trustworthy is enough to get people to be trustworthy in one-shot games and for partners to therefore expect trustworthy behavior and therefore trust. This is a possibility, one that does not present a difficulty for the theory developed later in this paper, but I do not believe that such a social contract would be understood by most to require trustworthy behavior in a contrived game any more than football players feel compelled to “be nice” when on the field.

20. By duty-based I mean that I have non-consequentialist or deontological notion of morality in mind in the sense that one will be trustworthy even if the harm done to any individual is essential zero.

21. Glaser, Laibson and Sacerdote (2002) are a notable exception.

22. A recent exception is Milyo (2004).

23. Labor is not human capital. Human capital is a stock, while labor is a flow. Human capital affects the nature of the flow of labor as a service.

24. Consider the similarity of networks to what goes on in business firms every day. Within many firms there is coordination among specialized divisions. This coordination is aided by repeated cooperative interactions and accumulated knowledge of how to work together. One reason the market value of the firm exceeds its accounting value is that networks within the firm make the firm more profitable.

Since the use of inter-firm networks does not use them up, it is properly regarded as a form of capital for the firm. If the firm suddenly ceased to exist, a loss would be incurred which would be the difference between the market value of the firm and the accounting value of its assets. Now consider a social network. If the embedded relationships that are a product of prior investments in social ties were to disappear overnight, the network would disappear, and this would lead to a reduction in the productivity of everyone who comprises the network (otherwise why be in the network?). Like capital generally, such a nexus of cooperative behavior (which is all a network is) has the property that its usage does not use it up (indeed, it likely strengthens with usage).

25. The sale of one's labor does not constitute alienability of human capital. Labor is a flow of services over time. Human capital is a stock that cannot be sold except under slavery. To be fair to Arrow, his comments were directly specifically toward physical capital.
26. Hirschman (1984) has argued that trust strengthens with use and decays with disuse [I got this from Dasgupta (2003, p. 6), the working paper version of a chapter that appeared in Ostrom and Ahn (2003). Elinor Ostrom (1999) makes a similar point about social capital per se.
27. Of course, the higher trust is upon initial contact, the more easily such networks will form and the more quickly they can be used to great effect. This is one reason why generalized trust is so important. It not only increases the scale of such social structures, it allows them to form more easily.
28. In many ways the idea of social capital reminds me of why we study macroeconomics – the idea that there are some phenomena that can only be understood if we pull back with a wide angle lens (example of explaining why an individual autoworker is out of work when it is ultimately the result of higher interest rates).
29. Indeed, in the neoclassical approach utility need not be measurable at all, even in principle, as the theory is now based on an ordinal concept of utility.
30. Teece (1992, p. 1) has stressed the importance of this balance: "Competition is essential to the innovation process and to capitalist economic development more

generally. But so is cooperation. The challenge ... is to find the right balance of competition and cooperation, and the appropriate institutional structures within which competition and cooperation ought to take place.”

31. Although some economists and social capital theorists have already argued (persuasively, in my opinion) that while individual social capital is human capital it is not merely human capital, they have not made this argument with respect to trustworthiness. Given the central role played by trustworthiness here and the unfamiliarity of some readers with these previous arguments, we take up this demonstration for trustworthiness now.
32. This might appear to contradict the idea of generalized trust. It does not in the following sense. Following Hardin (2002), think of trust as a 3-part relation: A trusts B to do X. Generalized trust is about A's willingness to trust B. Although such trust is not conditioned on who B is (unless B has been discovered by reputation to be untrustworthy), it is conditioned on X in conjunction with B. In other words, if X exposes A to a high level of potential exploitation, A might require more information before trusting B. In this sense, building-up a reputation for being very trustworthy makes others more willing to trust even in society with generalized trust in the sense that others will trust the individual even when temptation is likely to be strong.
33. Jean Ensminger is the best known proponent of this view. In her view, people are trusting when they believe others to be trustworthy. Others are trustworthy when institutions that sanction untrustworthy behavior are present and work well or in circumstances where repeated interactions are likely. This is consistent with what Hardin (2002) calls encapsulated trust, but its reach is very limited.
34. Social structures such as networks can be viewed as a kind of spontaneous order as defined by Hayek. Such structures can emerge from a critical mass of trustworthiness. This critical mass makes the spontaneous order a public good because the requirement of a critical mass before a qualitative change in the environment is achieved drives a wedge between the private returns to parental investments in trustworthiness and the social returns that result from a climate of trust.