Cyber-Trust

Jonathan, Cave and Lorenzo, Valeri

RAND Europe, University of Warwick

2004

Cyber-Trust

Jonathan CAVE[1], Lorenzo VALERI[2]
RAND Europe, 64 Maids Causeway, Cambridge, CB58DD, United Kingdom
[1]*Tel: +44 2476 523750, Fax: + 44 2476 523032, Email: cave@rand.org*
[2]*Tel: +44 1223 353329, Fax: +44 1223358845, Email: lvaleri@rand.org*

Abstract: Trust pervades the economic and societal interactions on which the Information Society is built. This paper applies economic analysis to issues of trust and identifies trust considerations affecting market structure, conduct and performance. The analysis highlights the impact of ISTs and the impact on a range of stakeholders.

1. Introduction

Trust is an essential attribute of electronic business and societal interactions, but the economic treatment of trust is still relatively simplistic. Trust pervades economics: it can be seen in the willingness to transact, use money, make contracts, invest, train, perform research and share intellectual property. Trust is also an economic 'good' in its own right – it can be produced, distributed, bought, sold and consumed. As a good, it can be viewed as social capital, a substitute for detailed formal contracts or extensive information or a complement to teamwork. Trust has externalities, and may even be viewed as a (local) public good. Trust is institutionalised in the use of reputations and is particularly important in information markets, where it resolves the 'paradox of information' (buyers cannot assess the value of information without access, and may then feel no obligation to pay for it or refrain from redistribution).

Trusting and trustworthy behaviour are learned responses – they may increase or decrease over time or make sudden jumps. External benefits and costs may lead to under- or over-supply of trust. More is not necessarily preferred to less, since trusting behaviour changes both the use of information and the taking of precautionary actions.

These generic characteristics are modified in electronic environments, and particularly in networked settings. Electronic transactions can support faster search and richer communication, but are vulnerable to new forms of opportunism. Moreover, virtual transactions may need higher levels of trust. In networked setting, the formation of links may be influenced by indirect network risks and through a 'market for certification.' Finally, trust can drive and reflect the formation of 'small worlds' for good and ill.

Policy that engages with economic activity on the Internet and the use of trust-enhancing strategies by members of the eEconomy require an appreciation of the reciprocal influence of trust and economic activity.

2. Objectives

If the institutions of the Information Society provide the warp of the social fabric, trust provides the weft – a softer complement giving resilience and support. Section 3 uses economics and game theory to analyse the evolution and impact of trust in electronic, networked environments and the design and valuation of trust-enhancing initiatives.

Complementing the application of economic analysis to the general issue of trust, Section 4 examines the role of trust in industrial economics.

A methodological objective, briefly summarised in Section 5, is the use of trust as a metaphor to analyse the evolution of social networks, particularly as related to ISTs. Trust influences behaviour, so we consider how coordination evolves in networks in particular, which norms emerge, whether global norms or local diversity will prevail and the influence of societal relationship structures. Link formation is itself an act of trust; a second strand considers underlying tensions between stability and efficiency in network formation – whether efficient networks will form and what can be done to encourage their formation. Further work [1] combines the two to examine the interaction of trusting/trustworthy behaviour and network formation – in particular, whether 'small worlds' are likely to form.

The final objective is to provide a framework for analysing the impacts of trust on various stakeholder groups to facilitate development and implementation of trust-improving strategies. The results should be of use to all those concerned with Information Society policy, including firms engaging in B2B and B2C eCommerce, eProcurement and eContracting, policymakers concerned with law enforcement, antitrust, industrial policy and eGovernment, and civil society groups concerned with public accountability and the societal consequences of engagement with the electronic polity.

3. Economic analysis of cyber-trust

Trust is a matter of expectation – extrapolation to other times and contingencies. Expectations may be bound up with process as well as outcomes. For example, one online customer may trust a transaction without distinguishing reliability of merchants, payment/delivery services, legal mechanisms that compensate losses, etc. Another with the same beliefs about the likelihood of different outcomes may evaluate purchases quite differently, being reluctant to disclose payment details to some (perhaps not all) agents. These considerations influence the nature and extent of participation, which in turn influence the effectiveness, the risks and the exposure of different parties.

Trust relationships are not necessarily symmetric, and it is useful to distinguish trusting from trustworthiness. Some aspects are summarised in Table 1.

Table 1: Trust and trusting

		Trusted Party		
		People	Systems	Organisations
Trusting party	People	Societal trust	Agency, privacy, accuracy	Reputation, Assurance
	Systems	Fault-tolerance	Complex system reliability	N/A
	Organisations	Agency	Reliance	Firm networking

From this, it follows that an appropriate distribution of trust (analogous to a distribution of risk) may be preferred to maximising trust (concentrating risk) as indicated in Table 2.

Table 2: Advantages of matching trust behaviour

	Trustworthy	Untrustworthy
Trusting	Appropriate delegation, specialisation	Enforcement costs, costs of adverse incidents
Untrusting	Excess contracting, monitoring costs; race-to-the-bottom.	Lost gains from trade, inappropriate risk allocation

This view of trust concentrates on trust as activity and overlooks the way societal institutions embody aspects of trust. An important literature [3] relates trust to contractual incompleteness – avoiding costly specification and monitoring of all contingencies and attached obligations. This applies to default contracts [2] and norms that allow markets to function; it lies at the heart of social capital [6][4][5]. This view is not uncontested. Some

see inherent conflict between formal contracts and trust; others note that trust in incomplete contracts involves acting on incomplete information, so "trust-enhancing" measures that provide assurance, indemnify against loss or add information weaken trust. Finally, the legal context of contract provides for monitoring, verification and enforcement in the event of breach and thus for trust hierarchies.

ISTs have at least three important impacts. First, they change the 'reach' of relational networks, enabling or forcing interactions with larger numbers of entities about whom less is known – e.g. the remote relationships of global ecommerce. They also change the 'bandwidth' of transactions that increasingly involve exchange of information that is at least commercially and possibly personally sensitive. The extension of trust is, as we argue below, unlikely to be monotone or smooth. Second, the institutions supporting trust are changed by the advent of new technologies. Some are weakened – for instance, greater anonymity may make reputations less effective, and globalised commerce may reduce the parties' ability to rely on common legal frameworks (or even to know which framework applies), and copyright and other IPR protection costs may outweigh their benefits. Other institutions may be enhanced – digital signatures can be firmly fixed to the parties, content and date and may be more secure than holographic ones. One likely consequence is emergence of public and private intermediaries to complement or replace existing institutions. Thirdly, the scale (possibly very small) and complexity (possibly large) of search and transaction in the Information Society have increased the use of artificial agents. Beyond important legal and practical considerations, it is not clear that trust extends smoothly to this world. We have seen that trust among people differs from trust between people and firms or governments. We must also consider whether people trust ICT systems: to act for them (agency); with information about themselves (confidentiality, privacy); to provide information they can safely act on (accuracy, currency, authentication, identity, integrity, etc.) – and whether interpersonal trust is enhanced by new e.g. surveillance) systems. The policy question is whether these impacts on trust work to reduce or improve the distribution of crime impacts and the burdens of crime reduction.

Formal models [1] support three main conclusions. First, there is a generic tension between efficiency and stability of trusting relationships – network formation equilibria typically involve "too much" or "too little" interconnection. This depends on how parties' interests are affected by linkage, and thus on ISTs and policies affecting their deployment. For instance, trusting exposes one to third parties trusted by the partner. Suppose ability to verify trustworthiness decreases as connections become less direct, and that direct links are costly to maintain. An efficient network maximises total payoff; any change from a Pareto efficient network makes at least one worse off. If costs are high, only the no-trust network is efficient; with intermediate costs star-shaped networks are efficient, and if costs are low the fully connected network is efficient. Trusted relationships are formed if all parties agree, and broken when at least one wishes. With intermediate costs, players only trust those with trusted relationships to others, but costs limit the number trusted to two. A network where each trusts exactly two others is a ring, but this cannot be stable - each could improve his payoff by breaking one link. Only the no-trust network is stable, but it is not even Pareto optimal - a line offers strictly more to both end and middle players.

Second, stable behaviour in a given network depends on ISTs and policy. To investigate this, we view trust as a symmetric 'coordination' game [10]. Table 3 shows a game with high-trust, low-trust and crime strategies. Players try to change strategy in response to their neighbours' choices. The unique stable outcome [11] in any network where everyone has the same number of links depends on the parameters as shown in Figure 1 below.

Table 3: Trust coordination game

	High Trust	Low Trust	Crime
High Trust	5, 5	*A*, *B*	-4, 3
Low Trust	*B*, *A*	3, 3	0, 0
Crime	3, -4	0, 0	1, 1

(assumes 5 > B and 3 > A see [1])

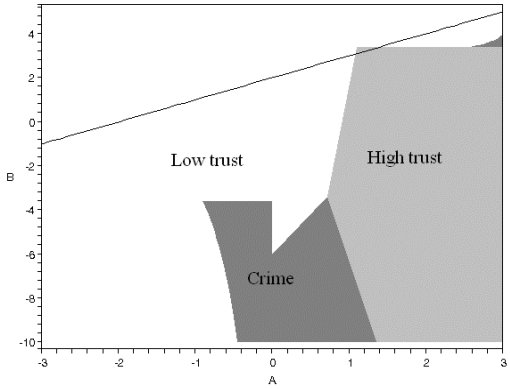

Figure 1: Stable conventions

Low-trust behaviour becomes more likely the payoff to trusting a low-trust person (A) falls relative to the low-trust partner's advantage of being trusted (B). Criminal behaviour prevails if B is low and there is relatively little difference between trusting and trying to cheat a low-trust partner (A between 0 and 3). Criminal behaviour reduces the scope for high-trust conventions – eliminating the criminal strategy makes high trust behaviour stable everywhere below the diagonal in Figure 1. Penalties for breach of trust or relative negligence liability standards – which reduce B - are less effective if criminal behaviour is possible – especially if victims of breach are not indemnified (A is low). If a low-trust person does not gain much from being trusted, insufficient indemnification may encourage criminal behaviour. Perhaps the main lesson is that policy should work with evolution to stabilise high trust (it is easer to reach the "High" zone than to raise A above 3). Also, trust is more likely where low trust involves precautions rather than selfishness (lower A).

Third, trust may evolve in uneven or perverse ways. Suppose networked individuals choose high- or low-trust channels based on expected net benefit and risk cost. Others do not observe individuals' weighting of these terms. Benefits combine a fixed component (e.g. lower transaction costs of certified public channels) and a 'network externality' varying with the number of high-trust partners. Under plausible assumptions, all sufficiently risk-averse players will avoid the high-trust channel. Figure 2 shows the equilibrium prevalence of high-trust individuals. When network effects or exposure are small the unique level of trust decreases with risk cost and increases with network effects. With high network effects there are high- and low-trust stable solutions decreasing in both risk cost and network effects and a perverse unstable middle-trust solution increasing in both parameters. If individuals learn about trust indirectly through news media or government reports the 'high trust' group is a representative sample of the population as a whole. This leads to the same general picture, but trust responds differently to policies or technologies that reduce perceived risk as shown in Table 4.

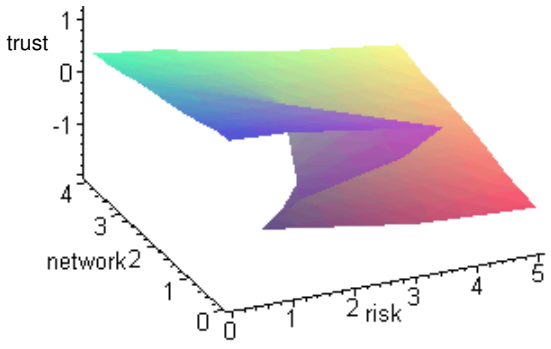

Figure 2: Trust, risk and network externalities

Table 4: Response of trust to fall in perceived risk

	Experience	Media
No network externalities	Fall	Constant
Peer-to-peer	Rise, then fall	Rise

(see [1])

4. Trust in cyber-economies

Trust is essential to commercial transactions when contractual completeness is costly or contract enforcement is unreliable. Trust among firms is traditionally associated with collusion, but increasingly also with informal 'networking' arrangements that reduce transactions and communications costs and improve collective efficiency. Trust between consumers and firms also saves costs and may provide incentives for competition and innovation leading to improved price and quality, but trust equally constitutes a barrier to customer switching (to the extent that it is relation-specific) and thus weakens competition.

4.1 Horizontal trust

Horizontal groups of firms trust each other to cope with an uncertain world or to counter specific challenges. Such trust may involve market allocation, price, quality or exchange of information. Firms wishing to agree strategies in pursuit of mutual interest face legal obstacles – hence the term 'trust' for cartels at the turn of the 19th century. To overcome this inability to contract and the combination of temptation and exposure, firms have developed trust-enhancing strategies: social contacts; interlocking directorates; most-favoured customer clauses, etc. Practices that facilitate collusion are 'trust based' if they are not directly enforceable or sanction activities that cannot be directly verified. Firms competing in electronic marketplaces have expanded opportunities to cloak departures from collusive agreements and a larger platform for their activities. This could increase the likelihood of defection – and thus the need to rely on trust. The same factors favour market-sharing agreements over price fixing. ISTs also enable rapid and effective responses to defection and facilitate extensive information exchange – enhancing firms' ability to trust each other. Finally, new forms of market contact reduce search and transactions costs, simultaneously raising the impact of defections on others, the likelihood of effective detection and the power, speed and credibility of retaliation.

4.2 Vertical trust

The impact of ISTs on trust and market power can be seen at search, payment, fulfilment and follow-up stages. New technologies can enhance search up to the point where a glut of rapidly changing information prevents effective comparison. This could lower the cost of differentiating products (inefficient monopolistic competition) or foster e.g. comparative search engines to sharpen price competition at a modest cost in deadweight loss. However, it is not obvious that search engines are entirely benign. They also prioritise comparative information, which may not serve consumers' interests or improve access to consumers by small-scale producers. If consumers regard firm size and persistence as signals trustworthiness, should they trust the information intermediaries responsible for that prominence (typically search engines assume no liability accuracy of information or quality of listed goods purchased on the basis of this information)? In particular, whether customers trust comparative search engines to identify good buys and whether such information increases customers' confidence that they have selected a best buy must be determined empirically. As yet the theory of vertical separation between sellers and information intermediaries is in its infancy [7].

Markets for information are even more prone to concentration than markets for 'real' goods and service. ISTs are more likely to enhance trust for standard commodities than for highly personalised or differentiated ones. Payment in the Information Economy typically involves financial intermediaries. Parties need to be assured of each other's identities, the reliability of payment and fulfilment and that repudiation risks are limited. The payer also needs assurance that financial information will not be 'reused' inappropriately. In contrast

to product search – where multiplicity and competition are the most trusted sources of assurance – trusted payment benefits from prominence and a degree of concentration. However, prominence brings risk – consumers worry about mistakes, abuse of trust and insufficient information security. Protection may weaken as financial service providers' reputation and market share increase, if these make them more attractive targets. ISTs affect fulfilment through globalisation of commerce – sellers may be located in other (even unknown) jurisdictions and consumer rights may be difficult or expensive to pursue. A dominant position may tie performance to reputation sufficiently to reinforce trust, but again there is a trade-off between reputational trust and monopoly power.

High quality can be signalled (and specific trust enhanced) by verified information (e.g. independent third party certification) or assurance. The efficacy of certification may be limited by psychological factors – in particular, whether information about risks heightens risk awareness. This is particularly true where information shifts effective or perceived risk allocation – decisions attract more liability the more information is provided in advance. Oligopolistic competitors may signal relative trustworthiness by talking up rivals' IST problems, but this may reduce trust in the market as a whole. Certification is an attractive alternative, but depends on reliability of the certifying authority. Much literature on cyber-notaries or Internet governance is concerned with the relative merits of competitive and coordinated certification; at the moment the key question is whether a public certifier of certifiers or an appropriate contingent liability scheme is the best way to answer "quid custodiet ipsos custodes?"

One final comment concerns the heavy information content of goods and services delivered over the Internet - buyer and seller cannot assess gains or divide from trade without an exchange of information, whereupon the buyer (or seller, in the case of personally identifiable information) can exploit the information without payment. A fair amount of trust is required to fit such transactions into the relatively anonymous framework of retail commerce. It may be hard to limit reuse without informational and contracting burdens that prevent some mutually beneficial transactions. This resembles the agency situation between patient and health care provider – effective treatment requires an exchange of information. The traditional approach to preventing abuse in this situation is one-sided liability – trust is tied to certification of health care providers and tort law. Even these arrangements have not survived technological advances, including availability of more effective treatments (which providers may wish to apply) and increased consumer information (due in no small part to health-related information on the Internet). ISTs' provision of information to the 'principal' (patient) undermines the former relationship of trust; patients no longer trust providers' expertise and providers no longer trust patients to follow advice or refrain from litigation.

4.3 Networked trust

The third type of relationship involves complements rather than substitutes or inputs. The IST world is increasingly a networked one, involving groups - trading partners, users of specific products, members of communities. Complementarity is the defining characteristic of such relationships [8]. A more sophisticated view distinguishes indirect from direct connections – here, the very act of joining a network involves trust that indirect connections will not damage his interests, and that partners will not to form damaging links. Consider the adoption of software services, which include some degree of security. Marginal costs of production are very low and the initial purchase decision may involve a commitment to further purchases of software products – often from other producers. The net present value of the supplier's customer base equals the customers' aggregate switching cost [9]. Incumbents try to maximise and potential entrants try to minimise switching costs. Both

strategies impose costs on the market as a whole. Moreover, 'churn' undermines trust in market stability and reduces supplier incentives to invest in durability and continuity. This further gives incumbent firms incentives to 'lock-in' suppliers of complementary products as a way of ultimately locking in consumers. In the specific context of security, this may mean lowering interoperability hurdles for 'compatible' products – effectively using the potential for attack at the interstices between software products as justification for extending the 'trusted zone' to enclose producers of complements.

Since these software products are used to mediate transactions and communication among networked people, the creation of proprietary standards builds a network externality among users of extended systems – the value to each person of using the system increases with the proportion of his or her contacts who use the same system. A dominant incumbent will attempt to maximise such network externalities. This leads to S-shaped adoption curves, 'tipping equilibrium' (capture of the whole market by a single 'extended standard') and, as in the trusted channel model in Figure 2, to abrupt jumps and local irreversibility leading to cyclic variation or dominance that survives radical technology changes.

In terms of market structure the first-mover advantage and need to capture suppliers of complements leads firms to reduce security barriers to developers, to share information with them and to shift the cost, complexity and liability burdens of security to customers. The implications for trust depend on whether customers can determine whether security provisions are effective (and appropriate). If customers cannot identify 'good' security, a form of Gresham's Law will operate - inferior security precautions will drive out good ones. Continuing demand must then be sustained by exaggerating threats. Finally, it should be noted that even effective security precautions may merely displace risk.

4.4 Consumer protection

Quality can also be signalled by assurance – warranty or compensation tied to breakdown of trust, analogous to penalties or liquidated damages. Two opposite strategies can signal quality and build trust. Extensive warranty credibly signals quality because it would be too expensive if quality were poor. Customers need not trust firms, but firms must trust customers not to make frivolous claims. A well-known firm can also credibly signal quality by minimal protection, placing its reputation on the line. This is informational trust - the consumer cannot control the risk transfer, merely whether to buy.

4.5 Liability

Where trust and crime are important there are externalities that can be mitigated by precautions. Civil, criminal and contract law allocate liability for consequences. If negotiation is costly, the allocation balances efficiency and fairness. Where externalities are one-sided (in other words, it is for me to trust you rather than for us to trust each other), so are the prevailing liability rules. Standard practice places liability on the least-cost avoider (or insurer). Where parties share blame or where avoidance costs are fairly symmetric, the rule is 'relative negligence'. An extreme case is criminal law - all parties bear full liability: this ensures efficient precaution, but is not efficient ex post. If parties can bargain costlessly over liability efficiency can be achieved regardless of how liability is assigned – but the eye-for-an-eye rule induces excessive precaution. These considerations seem particularly applicable to such risks as computer viruses, Spam, etc. Precautions have their own positive or negative externalities. The 'ownership' of risks and precautionary activities – and thus the trust placed in the system by participants (whether end consumers or B2B partners) – reflects monopolisation, the prevalence, adequacy and ownership of standards, and 'networking' among market participants.

5. Conclusions

Trust can reinforce IST tendencies towards market dominance. Competition to provide assurance can compensate for 'public-good' under provision of trust, stakeholders' natural disinclination to take appropriate precautions and inefficient risk allocation. The struggle to lock in information market customers can affect trust and trustworthiness.

Policy should recognise that more trust is not always better and exploit evolutionary forces surrounding trust. It should support self-regulation or appropriate liability trading. Other trust enhancement should be provided by public, not-for-profit, open bodies. Technology strongly affects trust but is liable to foreclosure, so policy should be framed in institutional rather than technological terms. Ownership of technical standards should carry a 'price' in terms of money, liability or responsibility – government's role may simply be to oversee markets on which prices are determined. Finally, sustainability of trust relationships may depend on asymmetry among the participants – 'improvements' that reduce asymmetry may actually undermine trust.

Trust is also a public good. Efficient provision depends on the incidence of costs and competence. Problematic possibilities include free riding and enclosure (creation of trust 'clubs' from which others are excluded).

Finally, analysis of trust as a societal norm shows that stability does not depend on whether trust is efficient, but rather on the balance of temptation and exposure. Society can be helped to evolve away from low trust lock-in at lower cost by evolution-aware policies. The speed of escape depends on clustering – cohesive self-referential groups can switch more rapidly than large diffuse groups and in so doing can 'seed' social transition. The results are significantly different when crime is included – policy interventions required to 'escape' the low trust outcome must be both more extensive and more precise, and policies that are "too weak" may undermine the rule of law and lead to 'criminal equilibrium.'

Space does not permit elaboration of applications in e.g. research networks, IPR, ecommerce, etc. The next steps are to develop the rendering of policies and validate the theoretical results empirically.

References

[1] Cave, J. (2004) "The Economics of Cybertrust" in R. Mansell (ed.) Cyber-trust and Crime Prevention, London: Elgar.

[2] Bacharach, M. and D. Gambetta (2001), 'Trust in Signs' in K. Cook (ed.) Trust in Society, New York: Russell Sage Foundation, pp. 148-84.

[3] Gintis, H. and S. Bowles (2004) "Persistent Parochialism: The Dynamics of Trust and Exclusion in Networks", Journal of Economic Behavior and Organization, forthcoming.

[4] Pollitt, M. (2001), "The Economics of Trust, Norms and Networks", Judge Institute Working Paper at http://www.econ.cam.ac.uk/electricity/people/pollitt/economicstrust.pdf.

[5] Puttnam, R. (2000), Bowling Alone - The Collapse and Revival of American Community, New York: Simon & Schuster.

[6] Fukuyama, F. (1995), Trust, New York: Free Press.

[7] Bailey, J. (1998), "Intermediation and Electronic Markets: Aggregation and Pricing in Internet Commerce", Unpublished PhD Dissertation, Program in Technology, Management and Policy, Massachusetts Institute of Technology.

[8] Katz, M. and Shapiro, C. (1994), "Systems Competition and Networks Effects", Journal of Economic Perspectives, 8: 93–115.

[9] Anderson, R. (2003), "Cryptography and Competition Policy: Issues with 'Trusted Computing'", Cambridge. Available at http://www.ftp.cl.cam.ac.uk/ftp/users/rja14/tcpa.pdf.

[10] Parties prefer to use compatible strategies. Ethical norms are inherently symmetric. Standard 'trust games' [2] do not have trust equilibria, but reputation extensions resemble asymmetric coordination games.

[11] For each norm, compute the minimum number of 'mistakes' needed to lead the population to evolve towards it from any alternative norm. The stable norm requires the fewest mistakes to displace a rival norm.