Perspectives

From Reciprocity to Reputation

This article is the third in a series of Perspectives on network capital and transaction costs. The first, "Richer Sourcing," was published in September 2004. The second, "Trusting Transactions," came out in March 2005.

Trust bypasses the rigidities of hierarchy and greases the wheels of markets. But it can take different forms. Technology is driving the substitution of one form of trust for another: *reputation* for *reciprocity*. This has consequences for strategy and for organization.

Reciprocity

We used to live in a world where information channels were expensive: they required proximity, trips, introductions, advertising, and so on. In such a "thin" network, the parties (whether corporations or people) are mutually vulnerable, depending on each other to convey true information and fulfill obligations. Apart from market contract or hierarchical control, we mediate this vulnerability through trust. The form of trust that we employ is reciprocity: we respond positively to cumulative, collaborative signals from the other party.

But the time and effort required to build trust through reciprocity are themselves a cost, further raising the cost of making those connections productive. When this feedback loop is strong enough, we economize on connections to the point where we operate within the thinnest possible trust network that still stays connected: the hierarchical org chart, the linear supply chain, or the loyal customer relationship, in which both the need and capacity for reciprocity are maximized while the number of channels is minimized. Thin, and also brittle, since these networks are all-or-nothing: we break a connection infrequently, but when we do, the break is sudden and complete and comes at a high cost.

Thin networks both require and enable thick relationships. Reciprocal trust can be transitive (A trusts B, and B trusts C, so A is willing to trust C), but such chains quickly peter out: reciprocity does not "scale." In general, reciprocal trusting relationships have to be built one at a time. Large swaths of organizational practice, marketing wisdom, and industry structure are based, unconsciously, on this logic.

Reciprocal trust of this kind is an asset, unrecognized by accounting conventions, and companies very often underinvest in it. The comparative innovativeness of Toyota's supply chain (as described in the first Perspective in this series) is in substantial measure explained by Toyota's reliance on (and willingness to invest in) reciprocal trust with its suppliers.

Reputation

But in a context where information channels are cheap, another kind of trust can flourish: trust based on reputation. A, with whom I do *not* have a trusting relationship, tells me about B. But the abundance of channels means that I can access lots of As. I may not trust any one of them (each one is a "noisy signal"), but I trust the aggregate: when there are *lots* of pieces of information, diffused over *lots* of routes, the noise (as information theorists tell us) simply cancels out. Moreover, B's behavior is affected by the knowledge that people are talking about her: she now has a *reputational asset* that she will rationally protect. As long as the gain from taking advantage of me is less than the consequent penalty to her reputation, I can be reasonably assured that she will not behave opportunistically, which gives me a further reason to trust her.

Reputation requires multiple and redundant information channels: ideally, everyone would communicate with everyone else. If this is assured, then reputation scales powerfully. In the language of information theory, double the size of a network and you multiply the signal-to-

Perspectives

FROM RECIPROCITY TO REPUTATION

noise ratio of a message everyone sends you by about the square root of two (this is the Metcalfe's law of reputation). In the language of common sense, the bigger the group, the greater the acuity, the more valuable (positive or negative) your reputation within it, and the more likely it is to influence your behavior.

And whereas reciprocity creates an asset that is specific to the reciprocating pair, reputation facilitates near-costless re-pairings among any of a network's members, since it is based not on *my* knowledge of you but on the *network's* knowledge of you (and my knowledge of your dependence on the network's knowledge). The trust asset belongs not to the pair but to the reputed individual or, in a different sense, to the network itself.

The catch, of course, is that in general there is no incentive, contractual or reciprocal, for the many As to give me information about B. Some positive incentive can be created with norms and by making one's own reputation a function of one's contributions to the reputations of others. But if communication imposes significant costs on the As, they will not proffer anything. The acuity and influence of reputation are constrained by the size and communication costs of the network within which it is embedded.

Technology

But technology changes all that, and therein lies the revolution. Technology enables cheap and redundant connections across which reputation can emerge. It enables parallel information channels ("trust technologies") for aggregating reputation: people (deliberately or implicitly) contribute information, and the technology summarizes the data, publishing the results at negligible cost. The blogosphere, for example, now comprises 27 million bloggers. Bloggers "vote" for each other through their blogrolls (fixed hyperlinks) and by citations. Sites such as Technorati provide navigation services and measure the "Authority" of bloggers on the basis of traffic patterns and citation.

Bloggers' ability to bestow reputation increases with the reputation they already have. Rather than the chaotic or random patterns one might expect, the resulting patterns of citation readership and Authority are quite ordered, bestowing on the top bloggers in each content domain extraordinary influence, extraordinary reputation. All without hierarchy, market transactions, reciprocity, or an ad budget.

As described in the second Perspective in this series, eBay's technology for aggregating ratings of buyers and sellers, Amazon's for creating reputations for books and for Amazon Marketplace vendors, and Google's for aggregating hyperlinks into votes on content all work in similar ways. But so too do Vault as a platform for employers' reputations; Slashdot, in which the karma of chatroom moderators is voted up or down by those who are moderated; del.icio.us, which aggregates 100,000 new bookmarks a day into a "folksonomy" of the Internet; and—the latest manifestation—virtual escort ratings within online games.

Edmunds.com (the site, not the company) probably has a bigger influence on car-buying decisions than the TV budget of the automotive industry. A recent Forrester study found that whereas about 5 percent of consumers trust telemarketing "somewhat" and nobody trusts it "completely," 50 percent trust consumer posts (as on Epinions) somewhat and a further 15 percent trust them completely. According to that crude metric, trust is 13 times higher in weak, unbiased signals than in strong, biased ones.

Moreover, we are at the beginning of all this. Reputation requires persistence of identity: I cannot build my reputation from anonymous transactions. And pseudonymous transactions in different domains build incommensurate reputations for my pseudonyms but not for me.

Perspectives

FROM RECIPROCITY TO REPUTATION

Technologies such as Trufina and Opinity already address those issues on a small scale, but in its next-generation Vista operating system, Microsoft will begin distributing an "identity metasystem" to hundreds of millions of desktops. This comprises a set of secure and open protocols that allow different authentication systems to work with one another. At first, the new protocols will make the current management of identity (logins, passwords, encryption, and so on) more user-friendly and much more secure. But very quickly the information covered by these technologies will broaden (under the control of the user). My "identity," if I choose, will become not just my username and password but also a vector of attributes (some legal, some descriptive, some preferential, some moral) that are variously relevant in various contexts, unambiguously tied to me, and certified by mechanisms as objective as a secure authentication from VeriSign and as subjective as the reviews of my last book.

With Vista and the technologies it enables, we will acquire the ability to aggregate and port this information costlessly and securely (but still under our own control) into any network we choose and to reveal it selectively to someone whose trust we need. We will have the ability to deploy certified information about ourselves in any way that builds the reputation(s) we desire.

The more people do this, and the more interconnected the human networks within which their reputations are embedded, the more valuable are those reputations, both in their information content (signal-to-noise) and—more important—as an earnest sign of their owners' moral reliability. Moreover, whatever the size of the population *not* participating in the reputational network, there is a self-sustaining process of positive selection: it will always pay for the "upper half" of nonparticipants to reveal their reputations, to achieve advantage over the "lower half," with whom they are otherwise being averaged. The bigger the reputation-sharing network, the greater the pressure to join: the classic increasing-returns pattern of network economics.

Implications

The implications of these developments are manifold. But consider just one. What if a company implemented these trust technologies internally? What if we all had reputations, aggregated from weak signals, that we could selectively reveal? What would that do to collaboration patterns? To the standard thin network/thick relationship paradigm of reciprocity? To the cost of wiring and (more important) *re*wiring the network? To the need for internal markets and hierarchy? To the need for the apparatus of human resources management? To the meaning of *leadership*? To behavior? To why people care?

And why stop there? Why not open the reputation network to people who are recruits and customers and suppliers? What would *that* do to behavior? And to brand? Would the company lose because the mixed signals are no longer "on message" as the recruiters and marketers and purchasing agents would define it? Or would the company—like Amazon with its book reviews—gain more from authenticity than it loses from honesty?

When it first emerged, the World Wide Web seemed a vast and alien toy, irrelevant to serious business. Now it is a daily tool. Well, now reputation-based networks such as the blogosphere have emerged: another vast but alien toy. Is the cycle about to repeat?

Philip Evans

Philip Evans is a senior vice president and director in the Boston office of The Boston Consulting Group.