Designing systems that people will trust

Patrick, Andrew; Marsh, Stephen; Briggs, P.

Security and Usability: Designing Secure Systems That People Can Use, 2005

National Research
Council Canada

Conseil national
de recherches Canada

Institute for
Information Technology

Institut de technologie
de l'information

NRC · CNRC

*Designing Systems that People Will Trust **

Patrick, A., Marsh, S., Briggs, P.
January 2005

* published in Security and Usability: Designing Secure Systems That People Can Use. January 25, 2005. NRC 47438.

Canada

00

Designing Systems That People Will Trust

Andrew S. Patrick, Pamela Briggs, and Stephen Marsh

Why Trust in Security

As we have seen in a previous chapter,[1] any security system is only as secure as its weakest link. Invariably, because of their social nature (and because of their human nature), the weakest links are often humans.[2] Thus, passwords get written on post-it notes and stuck to computer screens, or they become cycles of familiar words and numbers. In addition, social engineering succeeds in gaining inappropriate entry into supposedly secure systems because people will say things they're not supposed to, often to complete strangers who they just 'like.' Also, security systems are often turned off because they're too difficult to use, obscure, or downright impossible to comprehend for mere mortals.[3]

Because of the inherent social nature of computing systems, the issue of trust becomes paramount. It is not enough to design systems that are theoretically secure without taking into account the end users. Trust is a fundamental building block of society,[4] a means of making decisions about conferring authority or responsibility in unfamiliar or uncertain situations,[5] a method of understanding how decisions are made in context,[6] and one of the most important concepts in the security arena. Unfortunately, it also remains one of the most poorly understood concepts. A lack of trust will result in systems being ill-used at best, and not used at all at worst. A lack of understanding of trust, in both user and system, will result in the wrong decision being made in security contexts or no decision at all. Too much trust can be at least as dangerous as not enough, and not enough trust can be dangerous enough.

> A lack of trust will result in systems being ill-used at best, and not used at all at worst. A lack of understanding of trust, in both user and system, will result in the wrong decision being made in security contexts or no decision at all.

This chapter examines the issue of trust in security and privacy systems. These systems purportedly help users make decisions about who to trust with access, information, or data. For example, how much, when, and for what purposes can specific information be used? They can also help make decisions for the user when the user is not available. These decisions are based on a foundation of trust.

Current security systems are often seen as difficult to use, as getting in the user's way, or confusing the issue. As a result, they are often circumvented. A user should not have to delve into arcane issues of security to be able to allow access to a part of their personal information online: they don't have to in the real world, after all. In the real world, they rely on trust, an understanding of fiduciary responsibilities, and common sense. So it should be online.[7]

Fundamental questions arise when considering trust, including how to reliably represent trust in different interactions and interfaces, how to transform trust-based decisions into security decisions while maintaining the meaning of the trust-based decisions (in other words, attaining computational tractability without sacrificing meaning), how to transform in the opposite direction, and what the building blocks of trust really are in such contexts as information sharing or secure access to systems. Finally, because trust is fallible, what are its failings, how can they be addressed in this context, and what means of controlling the fallibility exist or should exist? Through investigating prior and current work in the area, this chapter arrives at recommendations for future systems and guidance for how they can be designed for use in a context of trust.

The rest of this chapter is organized as follows. In the next section we discuss the definitions of trust, and in the following section we examine the context of trust, its relation to risk, and the fundamental building blocks of trust online that have arisen from e-commerce research. Later, we present formal models of trust and describe what can be learned from these models. We conclude with a set of guidelines addressing how trust can be used in security systems, and concrete suggestions for system developers.

Definitions of Trust

Trust has not always been a subject of mainstream consideration.[8] In fact, prior to the Internet boom and bust, trust was a poor sibling to other sociological and psychological constructs. The Internet boom changed things as people began to realize that, with trust, people will buy things, and without it, they will not.[9] As simple as this may seem as an equation, it remains profound. What's more, the realization that imperfect designs can affect the trust of a user had an equally profound effect on how people went about implementing user interfaces, web sites, and interactivity in general.[10] The result has been an increasing amount of well-designed, well thought out interfaces, and a great deal of discussion in fields such as HCI and CSCW about how to encourage, maintain, and increase trust between people and machines, and between people and other people.[11]

Unfortunately, given all of this interest in trust, a deep and abiding problem arose: everyone knows what trust is, but no-one really knows how to define it to everyone's satisfaction. Thus, we end up with a great many different definitions, almost as many as there are papers on the subject, all of which bear some relation to each other, but have subtle differences that often cannot be reconciled. Trust, it seems, is a lot of things to a lot of people. Looking at the literature, this state of affairs is understandable because trust is multi-faceted, multi-dimensional, and not easy to tie down in a single space.[12]

The problem remains, however, that to discuss trust, one must in some way define terms. Recently trust researchers have backed away from this requirement since, when using the concept of trust in tools, it is often enough to know that trust *exists*, rather than what it *is*. Thus, we conjecture that, in cultures where the concept exists in some form, it is necessary only to acknowledge it in tools, and the user does the rest (and trusts, or not).

Whether or not this conjecture holds, definitions are at least for now expected. We suggest the following: "trust concerns a positive expectation regarding the behavior of somebody or something in a situation that entails risk to the trusting party."[13] In fact, this definition was arrived at after considerable study and reflection of trust research in numerous fields. The key here is that there is a judgment involved, as positive expectation, that there is free will on both sides to behave in certain ways, and that there is an element of risk. Problems remain with this and other definitions,[14] but it will do for our purposes.

> Trust can be defined as a positive expectation regarding the behavior of somebody or something in a situation that entails risk to the trusting party.

Given the multi-dimensional nature of trust, we have found it useful to discuss the different *layers* of trust, since it is these layers that affect how trust works in context. We have found that trust has three basic layers: *dispositional trust*, the psychological disposition or personality trait to be trusting or not; *learned trust*, a person's general tendency to trust, or not to trust, as a result of experience; and *situational trust*, in which basic tendencies are adjusted in response to situational cues.[15] These layers work together to produce sensible trusting behavior in situations which may or may not be familiar to the truster. For example, in an unfamiliar situation learned trust may be given less importance than dispositional trust (since no learned information is available), whereas a situation similar to others encountered in the past can allow a reliance on more learned trust. The situational trust allows cues, such as the amount of information or social expectations, act to adjust trust levels accordingly. Clearly, the more information available the better. Bear in mind, however, that a state of perfect information by definition removes the need to rely on trust.

Looked at in this manner, the goal of much HCI research and development is to create systems and interfaces that are as familiar as possible to the user such that the user need not make a (necessarily more limited) dispositional trusting decision, and to allow them to make a (more solid and comfortable) learned trusting decision. The goal of security and privacy systems is to allow the user to make these decisions with as many positive situational cues as possible, or to allow the user to provide and maintain their own situational cues in situations of less than perfect information, comfort, and ultimately, trust.

The Nature of Trust in the Digital Sphere

The concept of trust undergoes some interesting transformations when it is brought into the digital sphere. Whereas people may be quite adept at assessing the likely behavior of other people and the risks involved in the physical, face-to-face world, they may be less skilled when making judgments in online environments. For example, people may be too trusting online, perhaps routinely downloading software or having conversations in chat rooms without realizing the true behaviors of the other parties and the risks involved. People may also have too little trust in online situations, perhaps dogmatically avoiding e-commerce or e-government transactions in the belief that such actions cannot be done securely, at the cost of missed opportunities and added convenience.[16] Online users have to develop the knowledge needed to make good trust decisions, and developers must support them by making trustable designs.

One thing that is obvious is that trust in the digital sphere is negotiated differently from trust in face-to-face situations. Take the example of eBay—one of the most successful e-commerce businesses in operation today, and one in which complete strangers routinely send each other cheques in the mail (although this is becoming a less common means of payment as more sophisticated methods come on stream). How do eBay users develop sufficient trust in these unseen others to offset financial security concerns? Well, put simply, cues are provided in the form of a reputation system that not only enhances a sense of community among eBay members but also provides a profile of user experiences. These profiles are available to all vendors and customers—something that was unheard of in the world of offline commerce. Over years, the nature and utility of such cues has changed (and this will be discussed in more detail in a later section) but the principle that trust can be designed into a transaction was clearly established.

People are having difficulty knowing how to make trust decisions related to web security.

Another interesting example of the trust cues that can be provided to online users, and how difficult they can be to interpret, was provided in a study by Batya Friedman and her colleagues.[17] These researchers conducted detailed interviews of Internet users to explore the users' understanding of web security. They asked users to describe how they determine if a web connection is secure or not. The most frequent evidence was the appearance of the "https" protocol in the URL, and this was usually used correctly. On the other hand, the "lock" icon that appears in most browsers to indicate a secure connection was often misunderstood by the users, with many confusing the meaning of the open and closed locks. It was also common for people to use evidence about the point in the transaction (e.g., "this is the home page, so it is probably is not secure"), the type of information (e.g., "they are asking for my social security number, so it must be secure"), and the type of web site (e.g., "it is a bank, so they must be using security"). In addition, some people just made global mistrust decisions regardless of the evidence available (e.g., "I don't think any sites are secure"). This study makes it clear that people are having difficulty knowing how to make trust decisions related to web security.

"Phishing," the practice of creating mirror web sites of, for example, commerce or banking sites, and then sending emails to customers to ask them to "update their records urgently at the following [fake] link," is a particularly problematic exploitation of trust because it allows the fake site to obtain real account numbers, personal details, and passwords for subsequent fraudulent use on the real site. Phishing sites are often extremely sophisticated, often indistinguishable from the real site. Naturally, defenses against such attacks are difficult but possible. Some developers, for example, are creating web browser plug-ins that highlight the true location of a link, rather that the normal location display that can be easily obscured.[18] Ironically, recent features in web sites that are often seen as security concerns, such as using cookies to store login IDs and only asking for passwords, are an interesting defense—if I normally don't have to enter my ID, then a similar site that asks for the ID should be a clue about its authenticity.

Research is needed, however, on the effect of phishing attacks on trust in two domains. Firstly, it is clear that phishing, as with viruses and spam, directly affect the perception and feelings of trust in the Internet and online services in general. Secondly, and more interesting, it is unclear what effect a successful phishing attack has on the trust of the user who was compromised. That is, if I am a phishing victim through a fake bank site, how does this affect the amount of trust I have in the real bank site? The obvious answer is intuitive—that the brand is compromised. A more detailed answer should address how widespread the erosion of trust is, and what must be done to recover it. Phishing will likely remain a problem for some time, given the sophistication of the phishers. Strategies are necessary to combat the phenomenon. The Trusted Electronic Communications Forum[19] is one consortium aimed at tackling this and other spoofing attacks. Once again, social engineering is high on the agenda.

"Phishing" is a particularly problematic exploitation of trust.

Trust (and distrust) requires at least two parties: the truster and the trustee. It requires that the truster make an informed decision (that is, trust is not a subconscious choice, but requires thought, information, and an active truster). The converse is not true: it is not necessary for the trustee to know that the truster is, in fact, trusting them (it may be necessary for the trustee to know that *someone* trusts them, but that's a different debate). As discussed briefly above, in general it has been accepted that the trustee has to have some aspect of free will: that is, in this instance, they can do something that the truster would find untrustworthy. In the past this has been taken to mean that the trustee must be rational, conscious, and *real*: thus, machines cannot be *trusted*, they can only be *relied upon* (the difference is subtle, but not moot). In an age of autonomous agents, active web site, avatars, and increasingly complex systems, this argument is invalid. The corresponding argument that the trustee must know when he or she acts in an untrustworthy manner is somewhat more problematic. In any case, the phenomenon of anthropomorphism, whether validly directed or not, allows us to consider technologies as 'trustable' because people behave *as if* machines and technologies are trustable social entities that can in fact deceive us, and leave us feeling let down when trust is betrayed.[20]

The question remains then, especially when active entities such as autonomous agents or interactive interfaces are in mind, as to whom or what can *trust* and whom or what can be *trusted*. In this instance, one can consider humans as trusters and trustees, and computers in similar roles. Thus, we can consider trust between humans and humans, and between humans and computers, but we can also consider trust between computers and other computers, and finally, between computers and humans. Heretical as it may seem, it is the situations where computers are trusters (even as surrogate agents for humans) where the most power can be derived for privacy and security systems for several reasons. First, as was discussed above, humans are a weak link in the chain. Removing their social weaknesses from a security system can only strengthen the system against, for example, social engineering tactics. Second, since security and privacy decisions may have to be made in many circumstances at different times and in different places, removing the need for the human to make continuous decisions in this manner can only benefit the person. Third, there are circumstances where computer-based information needs to be moved around, and in these situations, computer to computer trust is not only desirable, but necessary.

In the circumstances where the truster is a computer, there is a need for a means by which the computer can 'think' about trust. Thus, a computationally tractable means of reasoning about trust is needed. It is not enough for the computer to be able to say "I trust you so I will share information with you." What information? How much? In which circumstance? In what context? We sometimes have a need to put some kind of value on trust (thus, "I trust you *this much*" is a much more powerful statement than "I trust you").

Of course, this leads to its own questions, such as what does "*this much*" actually mean, and how can we share trust and trust values. We address these questions below.

Formalisations and formal models of trust do exist and more are appearing regularly.[21] With each formalisation, old questions are answered, new questions arise, and we move closer to a real understanding of human trust and more capable trust-reasoning technologies. However, while formalizations exist, *computationally tractable* formalisations are much rarer. Unfortunately, it is these that are needed to better approach understanding, and to better approximate trusting behaviors in computers.

The Trust-Risk Relationship

As we have already argued, trust is a complex and highly context-dependent construct. In this section we address a key contextual issue in the relationship between trust and risk. This is an important issue because the models of trust captured in the literature (and described below) assume that trust is predicated on risk and that the processes underlying trust judgments vary as a function of the level of risk. This literature is largely concerned with trust in an e-commerce context, which is appropriate because privacy and security fears are salient for e-commerce, but we should note at the outset that the resulting models and recommendations are somewhat limited in addressing the broader security agenda. In particular, many of the issues relate to communications involving images and text which have as an outcome the purchase of a product (e.g., e-commerce on the WWW). Security systems for critical infrastructures outside of the Internet are likely to challenge our notions of trust in other ways—especially in terms of the computer-to-computer, agent-to-agent exchanges that are likely in the future.

The first point to note is that trust is intimately associated with risk—indeed it is possible to argue that in the absence of risk, trust is meaningless.[22] Let's take an everyday example: I could ask a stranger to look after my seat on a train (low risk) and not feel any need to engage in an evaluation of the trustworthiness of that stranger. However if I left an expensive video-camera or even my baby behind on the seat (high risk) then a more careful trust judgment would ensue. But this example raises other issues in relation to the trust-risk relationship. In particular it seems that the characteristics of trust are dependent upon the types of underlying risk. To pursue the example: if I would trust someone to watch my video-camera does that imply that I would trust them to look after my infant? Not necessarily because the two trust judgments are related but somehow distinct, with the latter relying more heavily on judgments of competence and kindness and the former on judgments of honesty. So to add to the argument made earlier, we may need to be able to phrase trust not just in terms of "I trust you this much" but also in terms of "I trust you *this much* to do *this thing*".

Trust is intimately associated with risk. In the absence of risk, trust is meaningless.

The same complexities occur in e-commerce. An online consumer's decision to trust an e-vendor may reflect beliefs about honesty, but is also likely to tap into decisions about competence and expertise, and it is further informed by judgments about the extent to which any information provided will remain private. Thus a seemingly simple act of trust invokes a complex set of judgments. Once again the risk assessment involved is crucial—there is no doubt that people are more willing to trust a site if perceived risk is low. This was shown very clearly in a study of over twenty-five hundred people who said they'd sought advice online.[23] Those that sought advice in relatively high risk domains (e.g., finance) were less likely to trust and subsequently act on the advice than those who sought advice in low risk domains (e.g., entertainment). Similar findings can be found in the well-known Cheskin/Sapient report on trust in e-commerce,[24] where, for lower-risk purchases such as books or groceries, trust was strongly associated with familiarity, whereas for high risk purchases, such as drugs or financial services, trust remained low, even when the companies themselves were well-known.

Even though some e-commerce transactions may seem to be low risk (say, involving small amounts of money) they usually involve high-risk elements such as the threat to privacy or credit card fraud. Furthermore, a typical exchange is complicated by uncertainties about whom or what is being trusted. Thus, in situations where perceived risk may be low, actual risks may be high and the assessment of actual risk is complex. For example, when a person logs into a secure web site to do a transaction, who are they trusting and what are they basing their trust decision on? In terms of people, they are trusting the writer of the web browser, the web host operator, the e-commerce vendor, all the network operators who handle their data, and the certificate authority that registered the web site.

Customers must be prepared to place their trust not only in the people, but also in the technology that underpins an interaction. Understanding the context for trust, therefore, involves understanding issues of encryption and data security as well as understanding the development of a psychological bond. Bollier, for example, argued that it is vital to distinguish between issues of "hard trust," involving authenticity, encryption, and security in transactions, and issues of "soft trust," involving human psychology, brand loyalty, and user-friendliness.[25]

Riegelsberger and Sasse have broken down the risks inherent in an e-commerce transaction in terms of firstly, risks that stem from the Internet, including: (a) whether credit card data gets intercepted, (b) whether the data is transmitted correctly and (c) whether the consumer uses the system correctly. Secondly, the these authors describe risks that are related to the physical absence of the online retailer, including (a) whether personal details will be kept confidential or transmitted to other parties and (b) whether the online vendor will actually deliver the products or services.[26]

People are faced with highly complex assessments of the risks they take when engaging in e-commerce transactions. One would assume that they would be influenced by the agencies charged with communicating information about the risk[27] and also the individuals or organizations charged with regulating the risk.[28] In e-commerce scenarios the regulation of security risk is usually the responsibility of the vendor, although trust is often gained by recourse to third-party endorsers offering seals of approval. However, consumers are surprisingly willing to accept risks when other trust indicators are present. Many Internet users will be familiar with a scenario in which they are asked to input detailed personal information about themselves in order to access the facilities available on a site. Users who input this information typically do so with the assumption that, first, the company honestly communicates its privacy policy and, second, that the company is capable of honoring those privacy claims. But few users actually spend the time checking this out, or even read the policies. In practice consumers seem to be more heavily influenced by the extent to which the facilities match their needs, whether the site has a professional look and feel, and also the extent to which the exchange seems predictable or familiar.[29] Indeed a very recent e-commerce study suggests that users are prepared to cast care to the wind and commit sensitive details to any site provided that the object of desire is compelling enough.[30] Once again, human fallibility can provide the weak link in the chain.

> A number of studies support a two-process model: people use cognitively intense analytical processing when the task is an important or particularly engaging one, whereas they use affect or other simple heuristics to guide their decisions when they lack the motivation or capacity to think properly about the issues involved.

Consumers are not always as cautious as they might be and it is possible to distinguish relatively 'hasty' and 'considered' processing strategies for the evaluation of trust in high and low risk environments. Chaiken identified two processing strategies by which an evaluation of trustworthiness may be made: firstly a heuristic strategy which follows a 'cognitive miser' principle—where people base decisions on only the most obvious or apparent information; and secondly, a systematic strategy that involves the detailed processing of message content.[31] Chaiken described two experiments that show that the degree of *involvement* in the issue affects the processing strategy. Those participants with low involvement adopted a heuristic approach to evaluating a message and were primarily influenced by the attractiveness, whereas those with high involvement adopted a systematic approach, presenting more arguments to support their judgment. A number of other studies in the persuasion literature support the two-process model, namely that people use cognitively intense analytical processing when the task is an important or particularly engaging one, whereas they use affect or other simple heuristics to guide their decisions when they lack the motivation or capacity to think properly about the issues involved.[32]

Such studies anticipate some recent findings with regard to online credibility. Stanford *et al*. invited experts and ordinary consumers to view health and finance information sites and found that experts (those having a high involvement with a site) were highly influenced by factors such as reputation, information quality and source, and perceived motive, in contrast to ordinary consumers (having a low involvement with the site) who were much more influenced by the attractiveness of site design.[33] The same is likely to be true of risk. In high-risk situations, or at least those situations that the user perceives as high risk, we would expect to see more evidence of careful analysis of trust indicators as opposed to a low-risk situation in which some rapid heuristic assumption of trust may be made. This high/low risk dichotomy is also played out in the trust literature where those experimental studies of initial trust where risk is imagined (would you buy from this web-site?) tend to place more emphasis on the attractiveness and the professional look and feel of sites, whereas those (few) studies which have actually involved substantive risk have emphasized careful consideration of integrity, credibility, and competence.[34]

It is worth saying something here about the relationship between trust and credibility. While a number of trust models incorporate judgments of source credibility in terms of expertise and reputation factors, and therefore see credibility as a component of trust, some researchers view trust as a component of credibility. Most notable is B.J. Fogg's work on the credibility of online information. Fogg is particularly concerned with the idea of the Internet as a persuasive technology. In a series of studies he and his colleagues at Stanford University have identified a number of factors that affect judgments of credibility. Positive factors included a real world feel to the site, ease of use, expertise, trustworthiness, and a site tailored to the individual. Negative factors included an overly commercial orientation and amateurism.[35] Fogg has interpreted this research in terms of a theory capable of explaining how web-credibility judgments are made. His prominence-interpretation theory posits two processes in the formation of a credibility judgment: prominence (the extent to which something is noticed) and interpretation (a considered judgment about the element under consideration). Fogg[36] argues that five factors affect prominence and three factors affect interpretation, as follows:

Prominence:

1. The involvement of the user in terms of their motivation and ability to scrutinize web content.
2. The topic of the web site
3. The nature of the user's task
4. The user's experience, and
5. Individual differences – for example in learning style or literacy level

Interpretation:

1. The assumptions in a user's mind (derived, from example, from cultural influences or past experiences)
2. The skills and knowledge a user brings to bear
3. The context for the user (in terms of environment, expectations, etc.)

There are interesting areas of overlap with the two-process model discussed earlier. Heuristic judgments clearly reflect the more 'prominent' aspects of an interaction, while analytic judgments reflect the interpretative processes outlined above. Perhaps the important issue for trust research is that the predictions made by prominence-interpretation theory (in terms of patterns of user-involvement, skills, and experience) are consistent with those derived from the two-process theory, and the guidelines that result are also in accord.

The Time-Course of Trust

The research on trust reviewed above suggests a need for more explicit consideration of the ways in which trust develops over time. It is certainly worth distinguishing between the kinds of trust that support transient interactions and those that support longer-term relationships.[37] A number of authors[38] have suggested that three phases are important: a phase of initial trust, followed by a more protracted exchange, which then may or may not lead to a longer-term trusting relationship. If one considers trust in this developmental context then some of the findings in the literature make more sense. In particular, consideration of a developmental context helps to reconcile the tension between those models of trust which suggest that it is a concept grounded in careful judgment of vendor expertise and experience, process predictability, degree of personalization and communication integrity,[39] and those models that suggest trust decisions depend much more heavily on the attractiveness and professional feel of a site.[40]

The importance of visual appeal in the early stages of interaction with a website is not unexpected given that in face-to-face interaction we often make judgments on the basis of the attractiveness of an individual, giving rise to the well known halo effect.[41] Other influences on first impressions in face-to-face conversation include the small-talk that strangers engage in. Some trust designers have tried to capture this in the design of relational agents that promote early trust. Thus Bickmore and Cassell describe the use of small talk to build 'like-mindedness' between interlocuters in the early stages of an interaction.[42] Although there is less documented research concerning trust in such interactions, the issue of how to make an agent trustworthy is likely to be important for future security systems.[43]

> The importance of visual appeal in the early stages of interaction with a website is not unexpected given that in face-to-face interaction we often make judgments on the basis of the attractiveness of an individual.

Another advantage of considering the developmental nature of trust is that it facilitates consideration of those factors that help to build trust and those that destroy it. A very early study of trust in automated systems demonstrated the intuitive finding that trust is slow to build up but can be destroyed very quickly.[44] This asymmetry is one of the reasons that researchers have suggested that the underlying processes involved in making or breaking trust are likely to be different. Thus, for example, McKnight *et al*.[45] describe two models, one for trust and one for distrust, and argue that disposition to trust and institution-based trust affects low/medium risk perceptions, while disposition to distrust and institution-based distrust will affect medium/high risk perceptions. The authors found that in contexts where people were merely exploring a site, the disposition to trust was most salient. Once they had made their mind up to engage in a higher-risk interaction with the site, the disposition to distrust became more important. McKnight *et al*. also found that promoting some initial exploration of the site was easy initially (because of the readiness to trust) and that this initial exploration could then be used subsequently to overcome the inclination to distrust when the user went on to engage in risky behavior. Interestingly, McKnight also observed a kind of halo effect such that a professional and well-designed site was associated with a disposition to trust.

These findings are consistent with the heuristic-systematic models described above if we consider that people are initially disinclined to look for hard evidence of trust (in the form of systematic assessment of expertise and careful investigation of privacy and security policies), but are instead happy to engage with sites on the basis that they are attractive and easy to use. This was certainly true of the users studied by Sillence *et al*.[46] in a study of trust and distrust of health sites. In their case a group of women searching for online information about the menopause and hormone replacement therapy were content to screen information first on the basis of attractiveness and design and subsequently went on to engage in more systematic processing to check the accuracy of the data and the motivation behind the site.

Models of Trust

Given the variety of definitions of trust, and the collection of components or dimensions that have been identified, it is not surprising that some people have attempted to simplify and concretize the domain by developed operationalizations and models. Researchers have developed a variety of models of trust components, antecedents, and/or consequences.[47] The advantage of models is that they may make fuzzy concepts clearer by defining terms and concepts. They can also provide structure where none existed before. More practically, developing a model may lead to specific metrics of interest that can be measured in research studies using questionnaires or other instruments. Models of trust can also lead to specific development advice. Some researchers working in the trust area, such as Egger, have used their models to develop criteria or checklists that practitioners can use to evaluate and improve a web site or similar service. In this section we review some of the models of trust, pointing out the similarities and differences, and we conclude with some specific lessons that the models can provide for developers.

> Models can make fuzzy concepts clearer by defining terms and concepts. They can also provide structure where none existed before. Developing a model may lead to specific metrics of interest that can be measured in research studies using questionnaires or other instruments. Models of trust can also lead to specific development advice.

Some of the earliest work on modeling trust focused on different components of the concept. Mayer *et al*.[48] proposed that trust is based on a set of beliefs about trustworthiness, and the most important beliefs concerned *ability*, *integrity*, and *benevolence*. Ability refers to the capacity for a trustee to be able to fulfill a promise made in a trusting relationship. Integrity relates to the promises made by the trustee—does he or she promise more than they can deliver. Benevolence refers to acting in another's best interest.

Gefen[49] operationalized this model of trust components by developing a questionnaire that addressed the three concepts of ability, integrity, and benevolence. Students who used the amazon.com website were asked questions related to Amazon's ability (e.g., "Amazon.com knows about books"), integrity (e.g., "I expect that Amazon.com will keep promises they make") and benevolence (e.g., "I expect that Amazon.com have good intentions toward me"). Analysis of the results showed that these concepts are reliable, statistically independent, and valid for predicting past shopping behavior and future intentions.

Bhattacherjee took a different approach and focused on the antecedents and consequences of trust for e-commerce situations.[50] That model consists of three components and, like many others, Bhattacherjee uses a flow diagram to illustrate the proposed relationship between the components, as is shown in Figure 00-1. The component of *familiarity* is defined as knowledge of the trustee based on prior interactions or experiences. *Trust* is assumed to be made up of beliefs in ability, benevolence, and integrity, based again on the pioneering work of Mayer *et al*. In this model, familiarity can lead to trust, which in turn can lead to a *willingness to transact*. In addition, familiarity can lead to a willingness to transact directly, even without feelings of trust. Such a situation might occur if a customer continues to transact with a vendor out of habit or convenience, even though there may be a lack of trust. Like Gefen, Bhattacherjee developed questionnaire items to operationalize each of the components in the model, and then demonstrated in an empirical study that the concepts were related in the expected statistical manner.

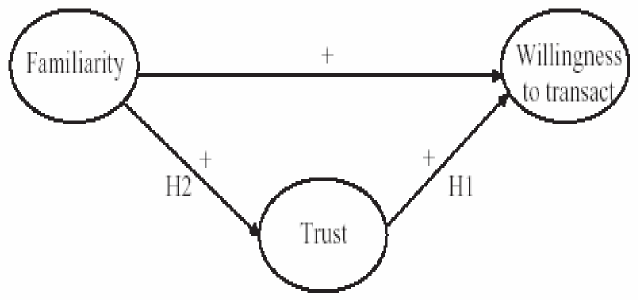

Figure 00-1. Bhattacherjee's Model of Trust

A similar model of trust for e-commerce was developed by Lee, Kim and Moon.[51] The model (see Figure 00-2) also describes antecedents to trust, this time focusing on three concepts: *comprehensive information*, *shared values*, and *communication*. In a way, these antecedents are describing the things that might be learned in the familiarity component proposed by Bhattacherjee, so the two models are similar in that respect. What makes the Lee *et al*. model unique is the addition of a *transaction cost* component that is seen as being in opposition to trust. In this model trust and cost are combined, in opposite directions, when customers make their decisions about e-commerce behaviors (in this case customer loyalty). Lee *et al*. describe three antecedents to transaction cost: *uncertainty*, the *number of competitors*, and *specificity* (the nature of the store or transaction). This model is important because it describes both trust and cost as being independent, opposing factors. According to the model, customers will choose to continue a relationship with a vendor if factors leading to trust are strong and factors leading to transaction costs are weak. We have recently adapted the model to replace transaction costs with the more general concept of perceived risk, and found it to be useful for explaining trust in a different domain.[52]

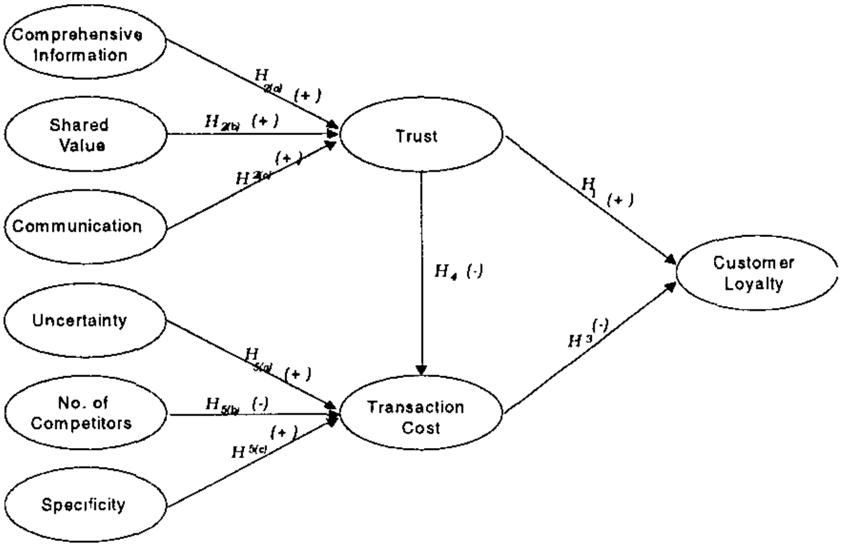

Figure 00-2. Model of Trust Proposed by Lee, Kim, and Moon

Corritore *et al*. also included trust and risk in their model (see Figure 00-3), although they proposed that increased perceptions of risk leads to decreased trust, rather than having trust and risk as independent factors.[53] This model also includes perceptions of *credibility* as a concept related to risk, and as we have seen, assessments of credibility are seen to be related to perceptions of honesty, expertise, predictability, and reputation. Corritore *et al*. also include *ease of use* in their model, and this is meant to measure how easy it is for a truster to achieve their goals (e.g., find the desired goods or complete the transaction). They propose that ease of use affects both credibility and perceptions of risk. Finally, this model also includes *external factors* that might affect a trust judgment. Such external factors include the environment or context of the transaction, the characteristics of the truster (e.g., a risk seeking or risk averse personality), the characteristics of the trustee (e.g., web site design), and the overall risk related to the transaction (e.g., the amount of money involved).

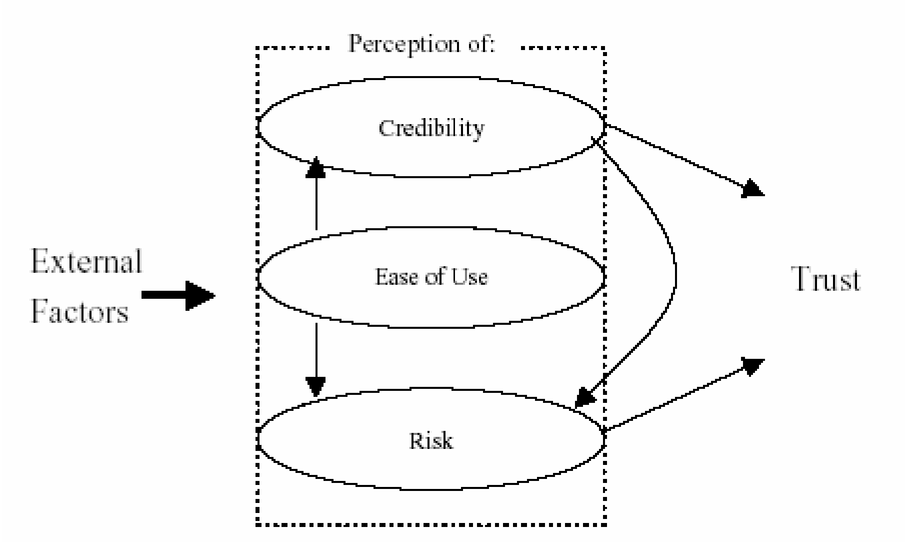

Figure 00-3. The Trust Model Proposed by Corritore et al.

In another model of trust in e-commerce situations, Egger also proposed an important role for external factors.[54] In Egger's MoTEC model (see Figure 00-4), *pre-interactional filters* are included to describe those factors in place before any transaction takes place. Included in this concept are factors such as the truster's disposition to trust, prior knowledge or experience, information and attitudes transferred from other (friends, the media, etc.), the reputation of the industry and company involved, and trust in information technologies and the Internet in general.

Two other important concepts in Egger's model are special roles for *interface properties* and *informational content*. Egger argues that interface properties, such as the visual appearance caused by graphic and visual designs, are important for creating first impressions that affect trust. Egger describes how trusters new to a situation or transaction make rapid assessments based on superficial cues about the usability, navigation, and reliability of an e-commerce application or web site. Later, the trusters may pay attention to the informational content in a slower, secondary phase of trust judgments. Here trusters assess competence and risk as they learn more about the transaction. Finally, MoTEC includes a *relationship management* component to explain the trust that may build up over time. Here trusters assess the responsiveness and helpfulness of a vendor, and how well transactions are completed over time, including fulfillment and after-sales support. In this way, Egger proposes a 3 stage model of trust: (1) rapid, superficial trust based on interface properties, (2) slower, reasoned trust based on an analysis of information content, and (3) relationship trust based on a history of transactions.

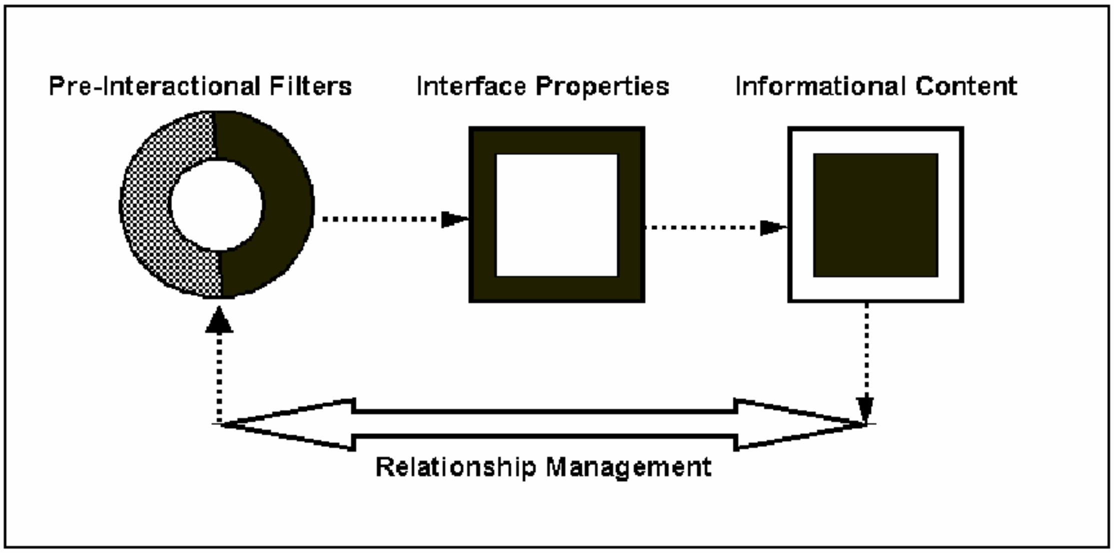

Figure 00-4. Egger's MoTEC Model of Trust

Taking into account the different models of trust described so far, each one proposing both common and distinct features or components, it is not surprising that there have been some attempts to build larger, more comprehensive models of trust. For example, McKnight and his colleagues[55] have developed a relatively complex model that includes many of the components proposed before (see Figure 00-5). This model outlines antecedent factors to trust in an e-commerce situation. Included are the factors of *disposition to trust* and trust in technology and the Internet (*institution-based trust*), as we have seen before. The model also includes the various attributes of a trustee, such as competence, integrity, and benevolence, which can contribute to *trusting beliefs*. One unique feature of the McKnight model is the distinction between *trusting intentions* and *trusting behaviors*. This is an important distinction because, although the theory of planned behavior[56] states that actions follow intentions, research shows that this is not always the case. It is one thing to state that you intend to do something (trust a vendor), but it may be quite another to actually do it. Very little trust research has actually measured true trusting behaviors, such as having trusters spend their own money in e-commerce transactions.

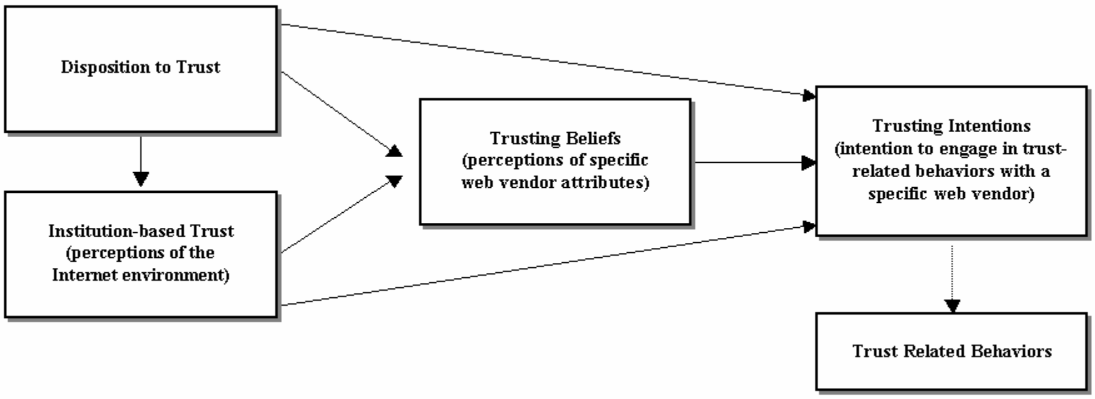
Disposition to Trust
Institution-based Trust (perceptions of the Internet environment)
Trusting Beliefs (perceptions of specific web vendor attributes)
Trusting Intentions (intention to engage in trust-related behaviors with a specific web vendor)
Trust Related Behaviors

Figure 00-5. The Web-Trust Model Proposed by McKnight et al.

McKnight and his colleagues have also made their concepts concrete by operationalizing them. Specific question topics for each concept are shown in Table 00-1, and questionnaires have been developed to ask trusters about each of these areas. As others have done, McKnight *et al*. have tested these questionnaires and shown that the concepts hold the statistic relationships that were predicted.

Table 00-1. McKnight et al.'s Operationalization of Trust Concepts

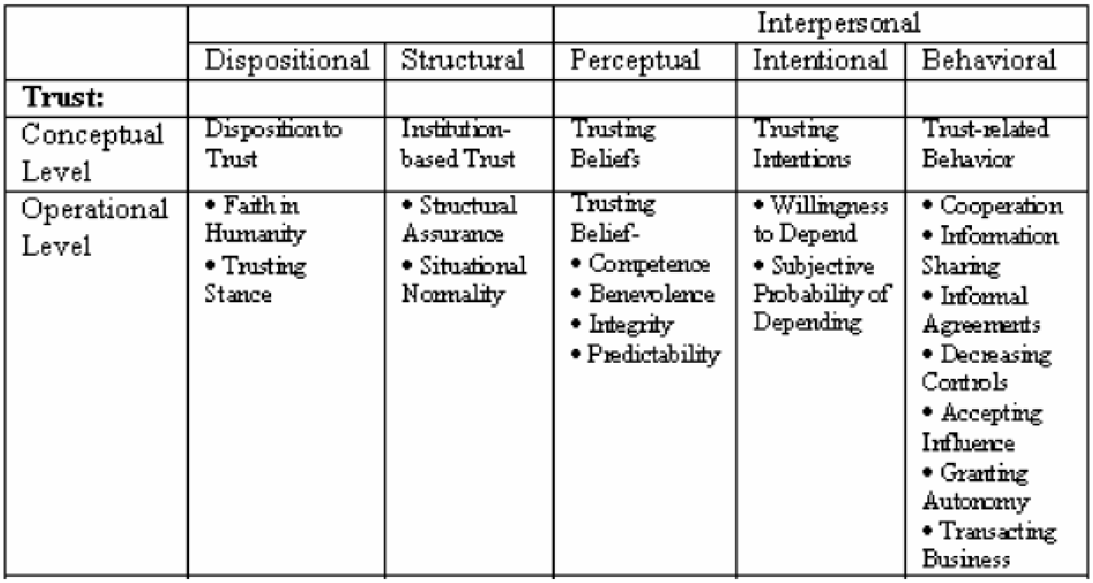

			Interpersonal		
	Dispositional	Structural	Perceptual	Intentional	Behavioral
Trust:					
Conceptual Level	Disposition to Trust	Institution-based Trust	Trusting Beliefs	Trusting Intentions	Trust-related Behavior
Operational Level	• Faith in Humanity • Trusting Stance	• Structural Assurance • Situational Normality	Trusting Belief- • Competence • Benevolence • Integrity • Predictability	• Willingness to Depend • Subjective Probability of Depending	• Cooperation • Information Sharing • Informal Agreements • Decreasing Controls • Accepting Influence • Granting Autonomy • Transacting Business

Another attempt at a comprehensive model has recently been described by Riegelsberger *et al*.[57] This model is somewhat different in that it focuses on the incentives for trustworthy behavior rather than opinions and beliefs about trust or perceptions of trustworthiness. This model (see Figure 00-6) describes the trust situation for both the truster and the trustee. Both the truster and trustee can choose to interact by performing *trusting actions* (truster) and *fulfilling promises* (trustee), or they can *withdraw* or *not fulfill*. Riegelsberger *et al*. describe what factors play a role in the decisions to take trusting actions and fulfill promises.

The first step in the model is for the actors to communicate by sending *signals* about a desire to interact. Often the situation is complex because the actors are *separated in space* (e.g., e-commerce buyers and sellers) and the actions are *separated in time* (e.g., delivery delays for e-commerce goods). Thus, signals can be important for showing trust-warranting properties. Signals are a method to demonstrate important intrinsic properties, such as benevolence, and they allow the actors to infer motivations and abilities.

Riegelsberger *et al*. also include contextual factors in their model (see Figure 00-7.), including temporal, social, and institutional properties. Social properties include things like reputation, while the institutional context is meant to convey things like the assurance given by job roles (e.g., bank tellers), regulations, and threats of punishment. This is similar to the "situation normality" concept included in the McKnight model. Riegelsberger *et al*. also describe different stages of trust that develop over time, and they discuss early, medium, and mature forms of trust. These concepts are not included in their model diagrams, however.

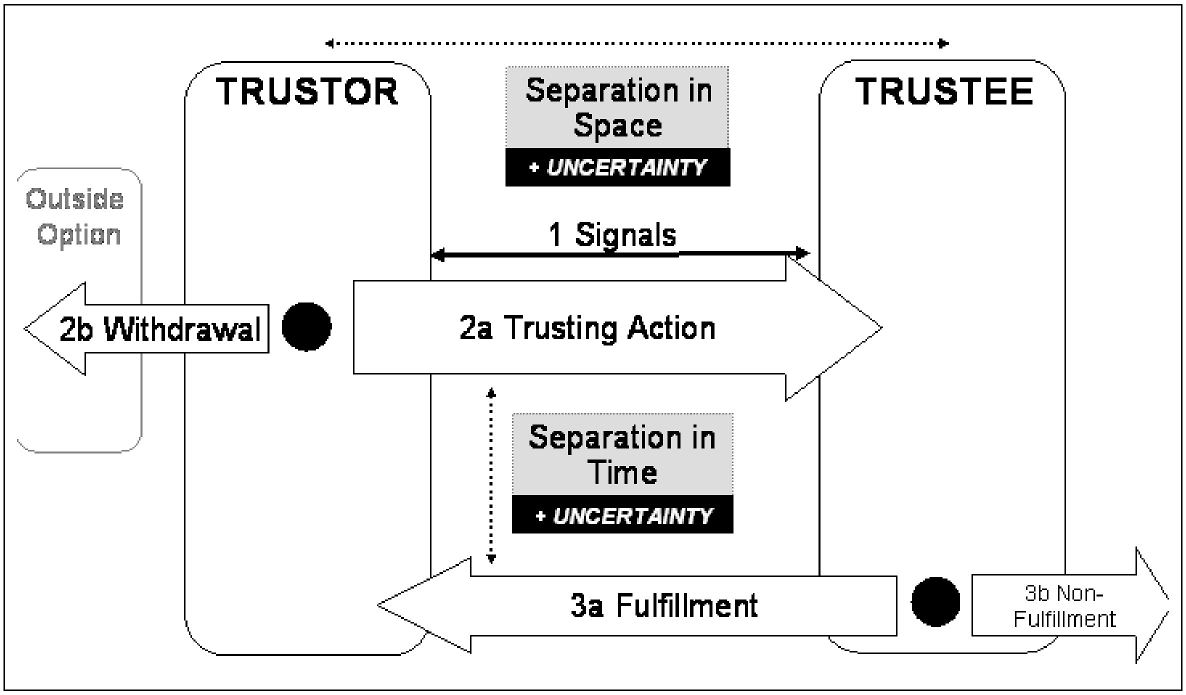

Figure 00-6. The Trust Situation Proposed by Riegelsberger et al.

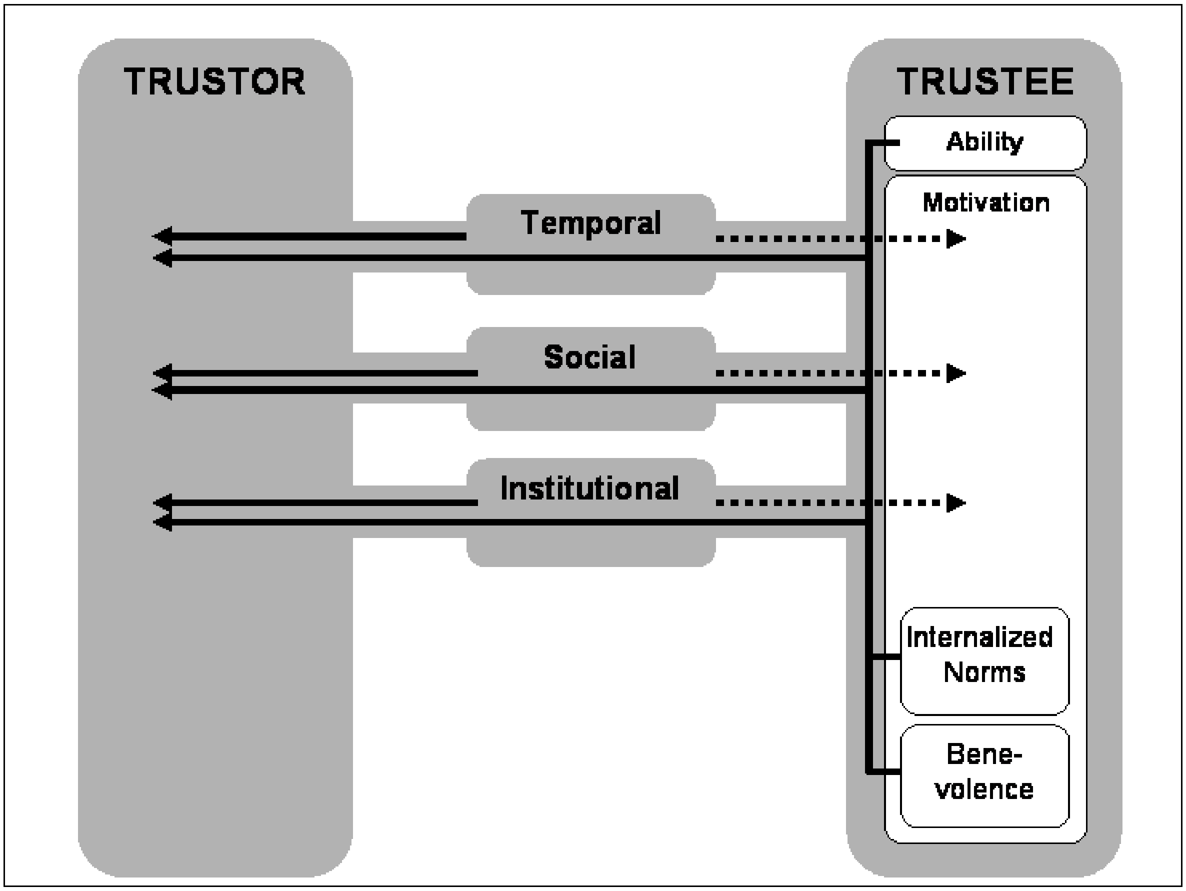

Figure 00-7. The Trust-Warranting Properties Proposed by Riegelsberger et al.

What are we to make of all these models? Although they all seem different, there are some common themes among them. More importantly, the research on the models can lead to specific advice to developers who wanted to build a trusted service.

- First, developers can learn that trust concepts can be operationalized into specific attributes or questions that can be examined in research and designs.
- Second, one of the key findings is that trust seems to be related be beliefs about another's ability, integrity, and benevolence.
- Third, trust and risk are related concepts, and factors that reduce risk perceptions, such as reducing uncertainty, can be beneficial for increasing trust or decreasing the need for trust.
- Fourth, ease-of-use characteristics, such as the ease of finding information and completing transactions, can affect trust.
- Fifth, external factors or context that may seem to be unrelated to the situation can affect trust, such as the characteristics of the truster and the type of risk involved in the transaction.
- Finally, trust probably develops in stages. In the first stage, superficial interface properties, such as colors and designs, can have a large effect on initial trust decisions. Later, users may make trust decisions based on more reasoned analysis of information. Eventually, long-term trust decisions are based on direct experience and personal service.

Trust Guidelines and Tools

The models developed above have not been produced in a vacuum. Almost without exception they have been developed with a view to making it easier for designers to identify those elements capable of promoting trust and those capable of destroying it. In this section we will briefly review some of the design recommendations that arise from this work and some of the tools that have been developed to aid the identification and promotion of trust factors in design. Once again the driver for much of this research is e-commerce and so many of the recommendations reflect specific issues to do with reassuring the online consumer. However, general issues of security accompany almost all e-commerce investigations and these guidelines are valuable to other situations.

A composite picture of trust guidelines (pulled from the literature reviewed above) is offered below – the order of the various guidelines suggests the point at which they are influential in interaction. Thus the lower numbered factors are likely to influence snap judgments made within seconds of visiting a site;[58] while the later factors are likely to come into play in the longer-term.

Trust Design Guidelines

1. Ensure good ease of use.
2. Use attractive design.
3. Create a professional image—avoiding spelling mistakes and other simple errors.
4. Don't mix advertising and content—avoid sales pitches and banner adverts.
5. Convey a 'real world' look and feel, for example with the use of high quality photographs of real places and people.
6. Maximise the consistency, familiarity, or predictability of an interaction both in terms of process and visually.
7. Include seals of approval such as TRUSTe.
8. Provide explanations, justifying the advice or information given.
9. Include independent peer evaluation such as references from past and current users and independent message boards.
10. Provide clearly stated security and privacy statements, and also rights to compensation and returns.
11. Include alternative views, including good links to independent sites within the same business area.
12. Include background information such as indicators of expertise and patterns of past performance.
13. Clearly assign responsibilities (to the vendor and the customer).
14. Ensure that communication remains open and responsive and offer order tracking or alternative means of getting in touch.
15. Offer a personalized service which takes account of each client's needs and preferences and reflects their social identity.

In addition to such guidelines, some authors have explicitly tried to develop tools to facilitate the consideration of trust factors in design and the subsequent evaluation of designs in terms of trust. One such is Florian Egger, whose MoTEC work was reviewed above. His model of trust for e-commerce led to the development of a trust toolbox consisting of three elements:

- GuideTEC: A set of design principles and guidelines relating to knowledge of the consumer and the vendor, properties of the interface, aspects of content, and the nature of the developing relationship with the consumer.
- CheckTEC: A 54 item checklist to be used in expert evaluations encompassing pre-interational filters, interface properties relating to appearance and usability, informational content relating to the company, products and services, security, privacy, and relationship management.
- QuoTEC: A 23 item questionnaire to elicit trust-related feedback directly from target customers.

Preliminary testing of these tools yielded positive results.[59] For example, checklist evaluators found 90% of the problems observed in user tests, whereas unguided evaluators found only 50%. The checklist is also interesting with respect to the security and privacy issues that are highlighted. The specific checklist items are detailed in the side bar below.

CheckTEC Checklist Items per MoTEC Dimensions

from Egger, 2003 (pg. 57):

1. Pre-interactional filters

1.1 This industry which this company belongs to is reputable

1.2 This company is known from the offline world or from advertisements

2. Interface Properties

2.1 Branding

2.1.1. The purpose of the website is clear from the start

2.1.2. The graphic design of the website is professional

2.1.3. The colour scheme and graphical elements are appropriate for this kind of website

2.1.4. The homepage incites users to explore the site further

2.1.5. The site pays attention to details, be they graphic, textual or navigational

2.1.6. Good use of grammar and spelling can be found throughout the site

2.1.7. The tone used in the texts is appropriate

2.1.8. The site is up to date

2.2 Usability

2.2.1. The pages display correctly in the most popular browsers

2.2.2. Legibility is high thanks to appropriate font sizes and contrast

2.2.3. The website is structured logically

2.2.4. Navigation across different sections of the site is consistent

2.2.5. Finding relevant information is made easy

2.2.6. The site contains no broken hyperlinks

2.2.7. It is easy to select items to purchase

2.2.8. It is easy to access the shopping basket and view its contents

2.2.9. It is easy to edit items in the shopping basket

2.2.10. The checkout and ordering process is intuitive

2.2.11. Appropriate feedback is given about the different steps in the transaction process

3. Informational Content

3.1. Company

3.1.1. The site provides complete offline contact details: e.g. physical address, phone and fax numbers,etc.

3.1.2. The site contains detailed information about the company's background

3.1.3. The site shows that there are real people behind the company: e.g. it contains key names, photographs and/or short biographies.

3.1.4. The site contains information about the company’s legal status, e.g. registration with a Chamber of Commerce

3.1.5. The company mentions partners involved in manufacturing, complementing or shipping products to show that it is reliable

3.1.6. The site contains precise information about when ordered items will be delivered

3.1.7. The site contains meaningful figures about the size of its customer base

3.2. Products & Services

3.2.1. Product descriptions are detailed and complete

3.2.2. Pictures (or other multimedia supports) effectively complement the textual descriptions

3.2.3. Product descriptions are objective

3.2.4. Sponsored content and advertisements are clearly labelled as such

3.2.5. All costs are displayed prominently and early in the transaction process

3.2.6. The site explains why some prices are unusually high or low

3.2.7. The site contains precise information about when orders will be delivered

3.2.8. The products/services sold on this website are familiar or from reputed brands

3.3. Security

3.3.1. The site features a prominent link to the security policy

3.3.2. The security policy clearly describes the measures taken to ensure that data is transferred, processed and stored securely.

3.3.3. The ordering process takes place on secure pages

3.3.4. When on a secure page, browser feedback is complemented with informative text and/or icons on the page itself

3.3.5. The site proposes alternative payment methods (i.e. not only credit cards)

3.3.6. The site features a detailed return policy

3.3.7. The site contains information about consumer redress mechanisms or financial compensation in case of fraud

3.3.8. The site contains a seal from a trusted third party that guarantees the company’s commitment to security

3.4. Privacy

3.4.1. The site features a prominent link to its privacy policy

3.4.2. The privacy policy clearly states what personal information is collected, how it will be used within the company and whether it will be sold to other companies

3.4.3. The site features a seal from a trusted third party that audits the company's privacy practices

3.4.4. The need for registration is delayed as long as possible

3.4.5. The site offers a clear overview of the information required in the registration or ordering form

3.4.6. Only personal information that is absolutely necessary is asked for in the registration or ordering form

4. Relationship Management

4.1. The site provides easily accessible online contact possibilities, e.g. e-mail addresses, live chat option, etc.

4.2. The site has a dedicated customer service area with frequently asked questions (FAQs) or a help section

4.3. The site makes it possible to manage and track orders online

Such guidelines (and related evaluation questions for the consumer) begin to suggest a way forward in designing trust for security systems, although they touch only the tip of the iceberg in terms of the research required. The time is now right for a proper agenda for trust research with specific respect to security and privacy systems, as opposed to only e-commerce.

Trust Designs

There are some examples available of successful designs that have promoted trust in online users. For example, gambling over the Internet using an off-shore, unregulated casino is an act that requires a great deal of trust. Such sites require that the gambler trust the casino operator to provide fair odds and handle money securely and properly. Shelat and Egger examined factors that online gamblers use when deciding to trust Internet gambling sites.[60] Conducted within the framework of the MoTEC model, the study revealed that informational content was the most important factor. People were most trusting when they could easily find information about the casino, its staff, and its policies. The second most important factor was relationship management, and trust-building attributes were an ability to contact the casino and rapid, high quality responses and payments. The third most important factor was interface properties, and this included usability and the ease of finding information. Finally, pre-interactional factors were the least important, with a positive attitude towards gambling being the most important determinant of trust in this category.

As stated earlier, one of the greatest success stories in terms of designing trust into a system is eBay. A number of trust design factors have been identified by Boyd[61] and they include (1) the use of a simple reputation system in which buyers and sellers give feedback about each other about issues such as promptness of payment; (2) the use of bulletin boards to reinforce the sense of community and to police undesired behavior and (3) a clear status system which relates not only to feedback but also to longevity with the vendor. This is reinforced with the use of icons such as the prestigious ‘shooting star’—an icon posted next to the usernames of people with a feedback rating of more than 10,000. Reputation systems are in operation in many sites, but Boyd notes that such design elements are cleverly worked into the community elements of eBay to reinforce the sense that its members genuinely help to build the company and are part of an “in group” of people engaged in an exciting venture.

Another example of a trusted design is the study that investigated the factors that lead to trust in online health advice.[62] This study examined the design factors that led a group of menopausal women to place their trust in sites that offered advice regarding hormone replacement therapy (HRT). The researchers found that most of the women preferred sites that were run by reputable organizations or had a medical or expert feel about them. They trusted the information on such websites especially when the credentials of the site and its authors were made explicit. Sites that indicated that the advice originated from a similar individual were also well received. Most participants showed some distrust of the advice and information on websites sponsored by pharmaceutical companies or those explicitly selling products. One of the most trusted sites was "Project Aware"—a "website by women for women." This site is split into menopause stage-specific areas, covers a wide variety of relevant topics, and provides links to original research materials. The language is clear and simple, and the layout is easy on the eye. Most importantly, however, the site establishes clear social identity signals, similar to those described for eBay, that tell readers that they are a member of a community and part of the in-group.

One point worth making about successful trust designs, however, is that they are only as trustworthy as the people that use them and trusted people can fail to be trustworthy, particularly when interacting with supposedly secure systems.[63] Trust design features do not in themselves guarantee a trustworthy system, and no amount of design work can compensate for a careless or malicious user. The phishing examples described at the start of this chapter provide food for thought—these attacks capitalize on our willingness to trust messages adorned with familiar and seemingly secure logos. Orgill *et al.* describe such 'social engineering' attacks and argue that ultimately user education will provide the best defense.[64] Certainly few users seem to fully evaluate the trustworthiness of different systems, even though they are influenced by the design factors described above.

Future Research Directions

We began this chapter with a discussion of some of the reasons why considerations of trust will be important for future privacy and security systems. Let us end the chapter with some explicit considerations of the trust issues raised by future technologies. We know that researchers and developers are increasingly excited about the concept of Ambient Intelligence (AmI). This term, first coined by the Advisory Group to the European Community's Information Society Technology Programme (ISTAG), refers to the convergence of ubiquitous computing, ubiquitous communication, and interfaces that are both socially aware and capable of adapting to the needs and preferences of the user. It evokes a near future in which humans will be surrounded by 'always-on', unobtrusive, interconnected intelligent objects, few of which will bear any resemblance to the computing devices of today.

One of the particular challenges of AmI, which distinguishes it from many other developments, is that the user will be involved in huge numbers of moment-to-moment exchanges of personal data without explicitly sanctioning each transaction. Today we already carry around devices (mobile phones, personal digital assistants) that exchange personal information with other devices, but we initiate most exchanges ourselves. In the future, devices embedded in the environment, and potentially in the body, will use software agents to communicate seamlessly about any number of different things: our present state of health, our preferences for what to eat, our schedule, our credentials, our destination, our need for a taxi to get us there in 10 minutes. Agent technologies will be required to manage the flow of information, and a great deal of exciting technical work is ongoing in this field. But many privacy and security concerns remain unanswered. How might we instruct these agents about when, where and to whom certain intensely personal details can be released?

> One of the challenges of Ambient Intelligence is that the user will be involved in huge numbers of moment-to-moment exchanges of personal data without explicitly sanctioning each transaction.

We are involved in several new research projects that address these issues, and some things have become clear. First, user-engagement in such technologies is crucial if we are to ensure a future devoid of suspicion and paranoia, but most users don't understand the complex technologies at issue here, and so new research methods inviting proper participation are required. Second, it is not enough to simply ask people about trust, privacy, or security in the abstract because what people say and what they do are two different things. Third, our future will be one in which many decisions are taken on our behalf by trusted third parties, so a great deal more information is required about the prerequisites for trust in regulatory bodies and agents. As we've already noted, there is a great deal of information available concerning building and breaking trust in e-commerce, yet only a very sparse literature on the ways in which people come to trust third parties in a mediated exchange. A related issue concerns the transfer of trust from one agent to another and recommender systems provide some interesting insights into this issue, particularly concerning the kinds of networks that support the transfer of trust from one individual to another. Finally we need to know a great deal more about what happens following loss of trust. From what we know already, it seems as though loss of trust can be quite catastrophic in a one-to-one relationship, but how does it percolate throughout a network of agents, each with their own set of trust indices? Such questions will be crucial for the development of privacy and security systems that people can genuinely trust.

Biographies

Andrew S. Patrick is a Senior Scientist at the National Research Council of Canada and an Adjunct Research Professor at Carleton University. He is currently conducting research on the human factors of security systems, trust decisions in privacy and e-commerce contexts, and advanced collaboration environments. Dr. Patrick holds a Ph.D. in Cognitive Psychology from the University of Western Ontario. http://www.andrewpatrick.ca

Professor Pamela Briggs currently holds a Chair in Applied Cognitive Psychology and the position of Acting Dean in the School of Psychology and Sport Sciences at the University of Northumbria, Newcastle upon Tyne, U.K. She is also Director of the PACT Lab—a new research laboratory for the investigation of Psychological Aspects of Communication Technologies. Pam has worked as a consultant for multinational organizations and her most recent work on trust and privacy issues in computer-mediated communication is funded by the Economic and Social Research Council's E-Society initiative.

Steve Marsh is a Research Officer at in the National Research Council's Institute for Information Technology (NRC-IIT), and is based in Moncton and Fredericton, New Brunswick. He is the Research Lead for IIT's Privacy, Security and Trust initiative. His research interests include trust, HCI, socially adept technologies, artificial life, multi agent systems, social computers, complex adaptive systems, critical infrastructure interdependencies, and advanced collaborative environments. http://www.stephenmarsh.ca/

End Notes

[1] see Sasse & Flechais, this volume. See also D. Schweitzer, "How to Toughen the Weakest Link in the Security Chain," *ComputerWorld* (2003). [Available Online.] [Cited September 27, 2004]; Available from http://www.computerworld.com/securitytopics/security/story/0,10801,77360,00.html.

[2] for a classic example, see Kevin Mitnick and William Simon, *The Art of Deception: Controlling the Human Element of Security* (John Wiley & Sons Inc, 2003).

[3] Anne Adams and Martina Angela Sasse, "Users Are Not The Enemy: Why Users Compromise Computer Security Mechanisms and How to Take Remedial Measures," *Communications of the ACM 42* (1999): 41-46.

[4] see for example Sissela Bok, *Lying: Moral Choice in Public and Private Life* (New York: Pantheon Books, 1978); Niklas Luhmann, *Trust and Power* (Chichester: Wiley, 1979); Barbara Misztal, *Trust in Modern Societies: The Search for the Bases of Social Order* (Cambridge: Polity Press, 1996).

[5] Bernard Barber, *The Logic and Limits of Trust* (New Brunswick, NJ: Rutgers University Press, 1983).

[6] Stephen Marsh, "Formalizing Trust as a Computational Concept," PhD Thesis, University of Stirling, Scotland, 1994; Mark Dibben, *Exploring Interpersonal Trust in the Entrepreneurial Venture* (Hampshire: MacMillan, 2000).

[7] to continue this discussion, see Barber, Luhmann, and also see Helen Nissenbaum, "How Computer Systems Embody Values," *IEEE Computer* (2001): 118-120.

[8] see Misztal and Luhmann.

[9] Cheskin Research & Studio Archetype/Sapient, "eCommerce Trust Study" (1999) [Available Online]. http://www.cheskin.com/think/studies/ecomtrust.html; Cheskin Research, "Trust in the Wired Americas" (2000) [Available Online] http://www.cheskin.com/p/ar.asp?mlid=7&arid=12&art=0

[10] Jakob Nielsen, "Trust or Bust: Communicating Trustworthiness in Web Design," *AlertBox* (1999). [Available Online.] http://www.useit.com/alertbox/990307.html

[11] see, for example Cheskin Research "eCommerce Trust Study" and Cheskin Research, "Trust in the Wired Americas." See also Ben Shneiderman, "Designing Trust into Online Experiences," *Communications of the ACM 43:12* (2000): 57-59; Gary Olson and Judith Olson, "Distance Matters," *Human–Computer Interaction 15* (2000): 139–178; Ye Diana Wang and Henry H. Emurian, "An Overview of Online Trust: Concepts, Elements, And Implications," *Computers in Human Behavior* (2005): 105-125; Cynthia L.Corritore, Beverly Kracher and Susan Wiedenbeck., "On-Line Trust: Concepts, Evolving Themes, a Model," *International Journal of Human-Computer Studies 58* (2003): 737-758; Jens M. Riegelsberger, Angela Sasse and John D. McCarthy, "The Researcher's Dilemma: Evaluating Trust in Computer-Mediated Communication," *International Journal of Human-Computer Studies 58* (2003): 759-781.

[12] Stephen Marsh and Mark Dibben, "The Role of Trust in Information Science and Technology," In B. Cronin (Ed.), *Annual Review of Information Science and Technology, 37* (2003): 465-498.

[13] Marsh and Dibben, 2003, p.470

[14] R.C. Mayer, J.H. Davis and F.D. Schoorman, "An Integrative Model of Organizational Trust," *Academy of Management Review 20:3* (1995): 709–734.

[15] Marsh & Dibben, 2003

[16] see Batya Friedman, Peter H. Khan, Jr. and Daniel C. Howe, "Trust Online," *Communications of the ACM 43:12* (2000): 34-40.

[17] Batya Friedman, David Hurley, Daniel C. Howe, Edward Felten, Helen Nissenbaum, "Users' Conceptions of Web Security: A Comparative Study," in *CHI '02 Extended Abstracts on Human Factors in Computing Systems, 2002*, 746-747.

[18] e.g., http://www.corestreet.com/spoofstick/

[19] see http://www.tecf.org

[20] see Byron Reeves and Cliffor Nass, *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places* (Stanford: CSLI Publications, 1996); B.J. Fogg, *Persuasive Technology: Using Computers to Change What We Think and Do* (Morgan Kaufmann, 2002); Cristiano Castelfranchi, "Artificial Liars: Why Computers Will (Necessarily) Deceive Us and Each Other", *Ethics and Information Technology 2:2* (2000): 113-119.

[21] see below and Stephen Marsh, "Formalizing Trust as a Computational Concept"; Alfarez Abdul-Rahman and Stephen Hailes, "A Distributed Trust Model," In *Proceedings of the ACM New Security Paradigms Workshop '97, Cumbria, UK. September 1997*; Cristiano Castelfranchi and R. Falcone, "Principles of Trust for MAS: Cognitive Anatomy, Social Importance, and Quantification," *Proceedings of the 3rd International Conference on Multi Agent Systems, 1998*, 72; Jonathan Carter and Ali A. Ghorbani, "Towards a Formalization of Value-Centric Trust in Agent Societies," to appear in *Journal of Web Intelligence and Agent Systems 2:3* (2004): 167-184.

[22] Andrew Brien, "Professional Ethics and the Culture of Trust," *Journal of Business Ethics 17* (1998): 391-409.

[23] Pamela Briggs, Bryan Burford, Antonella De Angeli and Paula Lynch, "Trust in Online Advice," *Social Science Computer Review 20:3* (2002): 321-332.

[24] Cheskin Research, "eCommerce Trust Study."

[25] David Bollier, *The Future of Electronic Commerce, A Report of the Fourth Annual Aspen Institute Roundtable on Information Technology* (Aspen, Colorado: The Aspen Institute, 1996).

[26] Jens Riegelsberger and Martina Angela Sasse "Trustbuilders and Trustbusters: The Role of Trust Cues in Interfaces to E-Commerce Applications," *Proceedings of the 1st IFIP Conference On E-Commerce, E-Business, E-Government, Zurich, 2001*. [Available Online.] http://www.cs.ucl.ac.uk/staff/jriegels/trustbuilders_and_trustbusters.htm

[27] O. Renn and D. Levine, "Credibility and Trust in Risk Communication," In R. Kasperson and P. J. Stallen (Eds), *Communicating Risk to the Public* (Dordrecht: Kluwer Academic Publishers, 1991), 175-218.

[28] W. Poortinga, and N.F. Pidgeon, "Exploring the Dimensionality of Trust in Risk Regulation," *Risk Analysis 23:5* (2003): 961-972.

[29] Briggs *et al.*, "Trust in Online Advice."

[30] Kathy Dudek, Pamela Briggs and Gitte Lindegaard, "Small Objects of Desire and Their Impact on Trust in E-Commerce," (in preparation).

[31] Shelley Chaiken, "Heuristic Versus Systematic Information Processing and the Use of Source Versus Message Cues in Persuasion," *Journal of Personality and Social Psychology, 39* (1980): 752-766.

[32] see, for example G. L. Clore, N. Schwarz and M. Conway, "Affective Causes and Consequences of Social Information Processing," In Robert. S. Wyer & Thomas. K. Srull (Eds.) *Handbook of Social Cognition* (Hillsdale, NJ: Erlbaum, 1994): 323-417; D.J. McCallister, "Affect-Based and Cognition-Based Trust as Foundations for Interpersonal Co-Operation in Organisations," *Academy of Management Journal*

38 (1995): 24-59; R.E. Petty and D.T. Wegener, "The Elaboration Likelihood Model: Current Status and Controversies," In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (New York: Guilford Press, 1999): 41-72; D. Albarracin, and G.T. Kumkale, "Affect As Information in Persuasion: A Model of Affect Identification and Discounting," *Journal of Personality and Social Psychology 84:3* (2003): 453-469.

[33] Julianne Stanford, Ellen R. Tauber, B.J. Fogg and Leslie Marable, "Experts vs. Online Consumers: A Comparative Credibility Study of Health and Finance Web Sites," [Available Online.] Consumer Web Watch [Accessed November 19, 2002]; Available from http://www.consumerwebwatch.org/news/report3_credibilityresearch/slicedbread abstract.htm.

[34] B. Chong, Z. Yang and M. Wong, "Asymmetrical Impact of Trustworthiness Attributes on Trust, Perceived Value and Purchase Intention: A Conceptual Framework for Cross-Cultural Study on Consumer Perception of Online Auction," *In Proceedings of ICEC 2003*.

[35] B.J. Fogg *et al.*, "What Makes A Web Site Credible? A Report on a Large Quantitative Study." *Proceedings of ACM CHI 2001 Conference on Human Factors in Computing Systems, 2001*, 61-68.

[36] B.J. Fogg, "Prominence-Interpretation Theory: Explaining How People Assess Credibility Online," *Proceedings of ACM CHI 2003 Conference on Human Factors in Computing Systems, 2003*, 722-723.

[37] e.g. D. Meyerson, K.E. Weick and R.M. Kramer, "Swift Trust and Temporary Groups," In R.M. Kramer and T.R. Tyler (Eds.), *Trust in organizations: Frontiers of theory and research* (Thousand Oaks, CA: Sage Publications, 1996), 166-195.

[38] Elizabeth Sillence, Pam Briggs, Lesley Fishwick and Peter Harris, "Trust and Mistrust of Online Health Sites," *Proceedings of the 2004 Conference on Human factors in Computing Systems, 2004*, 663-670; Florian Egger, "From Interactions to Transactions: Designing the Trust Experience for Business-to-Consumer Electronic Commerce", PhD Thesis, Eindhoven University of Technology, The Netherlands, 2003. [Available Online] ecommUSE [Cited December, 2005] Available from http://www.ecommuse.com/research/publications/thesis.htm.

[39] e.g., A. Bhattacherjee, "Individual Trust in Online Firms: Scale Development and Initial Trust," *Journal of Management Information Systems 19:1* (2002): 213–243; J. Lee, J. Kim and J.Y. Moon, "What Makes Internet Users Visit Cyber Stores Again? Key Design Factors for Customer Loyalty," *In Proceedings of CHI '2000, 2000*, 305-312; D.H. McKnight, V. Choudhury and C. Kacmar, "Developing and Validating Trust Measures for E-Commerce: An Integrative Typology," *Information Systems Research 13:3* (2002): 334-359.

[40] U. Steinbruck, H. Schaumburg, S. Duda and T. Kreuger, "A Picture Says More Than a Thousand Words—Photographs As Trust Builders in E-Commerce Websites," *Proceedings of Conference on Human Factors in Computing Systems CHI 2002 (Extended Abstracts), 2002*, 748–749.

[41] see for example N.R. Bardack and F. T. McAndrew, "The Influence of Physical Attractiveness and Manner of Dress on Success in a Simulated Personnel Decision," *Journal of Social Psychology, 125* (1985): 777-778; K. Dion, E. Bersheid and E. Walster, "What is Beautiful is Good," *Journal of Personality and Social Psychology, 24* (1972): 285-290.

[42] T. Bickmore and J. Cassell, "Relational Agents: A Model and Implementation of Building User Trust," *Proceedings of Conference on Human Factors in Computing Systems CHI 2001, 2001*, 396-403.

[43] Andrew S. Patrick, "Building Trustworthy Software Agents," *IEEE Internet Computing 6:6* (2002): 46-53.

[44] J. Lee and N. Moray, "Trust, Control Strategies and Allocation of Function in Human–Machine Systems," *Ergonomics 35:10* (1992): 1243–1270.

[45] D. Harrison McKnight and Norman L. Chervany, “Trust and Distrust Definitions: One Bite at a Time,” In R. Falcone, M. Singh, and Y.-H. Tan (Eds.): *Trust in Cyber-societies, LNAI 2246*, (Springer, 2001) 27–54.

[46] Elizabeth Sillence *et al.* “Trust and Mistrust of Online Health Sites.”

[47] for a review, see Sonja Grabner-Krauter and Ewald A. Kaluscha, “Empirical Research in On-Line Trust: A Review and Critical Assessment,” *International Journal of Human-Computer Studies, 58* (2003): 783-812.

[48] R.C. Mayer, J.H. Davis and F.D. Schoorman, “An Integrative Model of Organizational Trust,” *Academy of Management Review 20:3* (1995): 709-734.

[49] D. Gefen, “Reflections on the Dimensions of Trust and Trustworthiness Among Online Consumers,” *The DATA BASE for Advances in Information Systems, 33:3* (2002): 38-53.

[50] A. Bhattacherjee, “Individual Trust in Online Firms: Scale Development and Initial Trust.”

[51] J. Lee *et al.*, “What Makes Internet Users Visit Cyber Stores Again? Key Design Factors for Customer Loyalty.”

[52] Andrew S. Patrick “Building Trustworthy Software Agents.”

[53] Cynthia L. Corritore *et al.*, “On-Line Trust: Concepts, Evolving Themes, a Model.”

[54] Florian Egger, “From Interactions to Transactions: Designing the Trust Experience for Business-to-Consumer Electronic Commerce.”

[55] McKnight *et al.*, “Developing and Validating Trust Measures for E-Commerce: An Integrative Typology,” McKnight, D.H, Chervany, N.L., “What Trust Means in E-Commerce Customer Relationships: An Interdisciplinary Conceptual Typology.”

[56] I. Ajzen, “The Theory of Planned Behavior,” *Organizational Behavior and Human Decision Processes, 50* (1991): 179–211.

[57] Jens Riegelsberger, M. Angelas Sasse and John D. McCarthy, “The Mechanics of Trust: A Framework for Research and Design,” *International Journal of Human Computer Studies*, in press.

[58] see Elizabeth Sillence *et al.*, “Trust and Mistrust of Online Health Sites.”

[59] Florian Egger, “From Interactions to Transactions: Designing the Trust Experience for Business-to-Consumer Electronic Commerce.”

[60] B. Shelat and Florian Egger “What Makes People Trust Online Gambling Sites?” In *Proceedings of the Conference on Human Factors in Computing Systems, 2002 Extended Abstracts*, 852-853.

[61] J. Boyd, “In Community we Trust: Online Security Communication at eBay,” *Journal of Computer-Mediated Communication 7:3* (2002) [Available Online] Available at: http://www.ascusc.org/jcmc/vol7/issue3/boyd.html.

[62] Elizabeth Sillence *et al.*, “Trust and Mistrust of Online Health Sites.”

[63] Gregory L. Orgill, Gordon W. Romney, Michael G. Bailey and Paul M. Orgill, “The Urgency for Effective User Privacy-Education to Counter Social Engineering Attacks on Secure Computer Systems,” *In Proceedings of the 5th Conference on Information Technology Education, 2004,* 177-181.

[64] Orgill *et al.*