

LOOKING WITHOUT SEEING: UNDERSTANDING UNSOPHISTICATED CONSUMERS' SUCCESS AND FAILURE TO DETECT INTERNET DECEPTION

Stefano Grazioli

Management Sciences and Information
Systems Department
University of Texas at Austin
grazioli@mail.utexas.edu

Alex Wang

Department of Advertising
University of Texas at Austin

Abstract

Do unsophisticated consumers fall prey to Internet consumer frauds? Why? To answer these questions this paper integrates two streams of empirical research: the process-oriented theory of deception, and the broader deception, trust, and risk (DTR) model of Internet consumer behavior. A laboratory experiment tests several alternative hypotheses about the determinants of failure at detecting Internet deceptions. The findings suggest that Internet consumers process the clues that a site may be deceptive, but are unable to effectively evaluate and combine these clues, i.e., to draw correct conclusions from them. This is good news in the ongoing struggle against Internet fraud because it suggests that consumers lack the knowledge, not the capacity, to detect deceptions and that consumer education programs might be effective in helping consumers to protect themselves.

Keywords: Internet fraud, Internet deception, trust, information security.

INTRODUCTION

Monitoring agencies such as the Federal Trade Commission, the Securities and Exchange Commission, and the Federal Bureau of Investigation have recently voiced concerns over the growing occurrence of Internet consumer fraud and have started specialized programs targeted at detecting and prosecuting it (e.g., Pitofsky 1998). In addition to financial damage, victims often suffer from the psychological discomfort of being victimized, the loss of time necessary to file complaints and refund requests, and the theft of private information. Recent research results suggest that the median loss per victim is more than \$700 (Grazioli and Jarvenpaa 2000).

While the technical literature on Internet security is rapidly growing, comparatively little is known about the nature of Internet fraud, its victims, or its perpetrators. This paper looks at the victims. It starts from the conjecture that not every individual is equally susceptible to deception (Ekman 1992) and that useful knowledge can be harvested by comparing those who manage to steer clear of malicious web sites with those who do not. What we learn can be used to educate the public and focus the resources of monitoring agencies.

The next section integrates two current streams of research on deception and proposes a model of how Internet consumers make up their minds about purchasing at an unfamiliar (no brand-name) web store. The third section describes a lab experiment with unsophisticated users designed to test several hypotheses about the determinants of success and failure at detecting Internet deception. The fourth and fifth sections discuss the results and conclusions.

INTERNET DECEPTION, TRUST, AND RISK

According to the theory of deception (Johnson et al. 1992, 1993, 2001; Mitchell and Thompson 1986; Thagard 1992), a deception is a cognitive interaction between two parties under conflict of interest. One party, the deceiver, manipulates the environment of the other party, the victim, so as to foster in the victim an *incorrect representation* of his/her situation and therefore bring about a desired action.

The Internet medium affects the nature of existing forms of deception because (1) it makes identity (of items of exchange, individuals, and organizations) easy to falsify and difficult to authenticate (Jarvenpaa and Grazioli 1999); (2) it has decreased the economic resources needed to set-up a credible-looking storefront; (3) it has provided deceivers with increased access to potential victims; and (4) it has made the proceeds of crime easier to secure not only anonymously but also in jurisdictions where pursuing perpetrators is difficult (Morris-Cotterill 1999).

Every time that an information-technology mediated phenomenon is investigated, it makes good scientific sense to ask: "What is new here?" We argue that the specifics of Internet technology have introduced the possibility of new forms of crime, some of which have already appeared "in the wild." This paper investigates one instance of such new forms: "page jacking" (FTC v. Pereira et. al. 1999; Jarvenpaa and Grazioli 1999). Page jacking consists of maliciously redirecting an Internet user from his intended target site to another location. Page jacking is especially pernicious when the malicious location is similar enough to the original so that the user does not realize that page jacking has happened and believes he is at the intended site. This incorrect representation of environment may lead the victims to engage in behavior in which they would not engage otherwise. According to the Federal Trade Commission, about 2% of the pages on the web have been targeted by this malicious practice (Pitofsky 1998).

This study integrates the deception, trust, and risk (DTR) model by Grazioli and Jarvenpaa and the theory of deception by Johnson and his colleagues. The DTR model is a broad model of the determinants of the decision to purchase on-line. The theory of deception focuses on how individuals detect deception in information-intensive environments. The DTR model assumes that consumers evaluate the deceptiveness of a web store as part of their overall purchase decision. One of the main weaknesses of the DTR model is that it does not tell us *how* the evaluation of deceptiveness happens. By contrast, the theory of deception provides us with a process model of how potential victims evaluate the deceptiveness of information provided to them, but does not explain how that determination fits in the overall purchase decision. By combining these two approaches, we can leverage their strengths.

The DTR Model

The DTR model explains Internet consumers' willingness to buy items on-line from unfamiliar sites. The model draws from the theory of reasoned action (Fishbein and Azjen 1975) in proposing that willingness to buy depends on the consumers' attitude toward the web store (see Figure 1). Attitude is defined as an either favorable or unfavorable evaluation of the site.

Following Kramer (1999), DTR defines **trust** as a state of perceived vulnerability that is derived from an individual's uncertainty regarding the motives, intentions, and prospective actions of others on whom they depend. DTR emphasizes that trust plays a key role in the determination of attitudes toward a web store (Crosby et al. 1990; Ganesan 1994; Macintosh and Lockshin 1997; McKnight et al. 1998). **Risk** refers to a consumer's perceptions of the uncertainty and adverse consequences of engaging in an activity (Dasgupta 1988; Dowling and Staelin 1994; Kelley and Thibaut 1978; Kramer 1999).

The model hypothesizes that willingness to purchase (or to recommend a purchase) from a web store depends on the consumer's attitude toward the store (hypothesis H1). The consumer's attitude is determined by trust (hypothesis H2a) and perceived risk (H2b). In addition, trust moderates the relationship between risk and attitudes (H2c). When trust is high, risk considerations have less of an impact on the formation of attitudes about the site (Das and Teng 1998; Kollock 1994).

Trust and the perception of risk are influenced by two sets of mechanisms through which consumers assess risk and trust: assurance mechanisms and trust-building mechanisms.

Assurance mechanisms deal with conditions that reduce the probability of deceitful behavior or increase the penalty for detected opportunistic behavior (Spreitzer and Mishra 1999). DTR identifies four specific types of assurance mechanisms: (1) third-party seals (e.g., the Better Business Bureau "BBB On-Line" seal), (2) warranties, (3) news clips from third-party publications, and (4) physical store location.

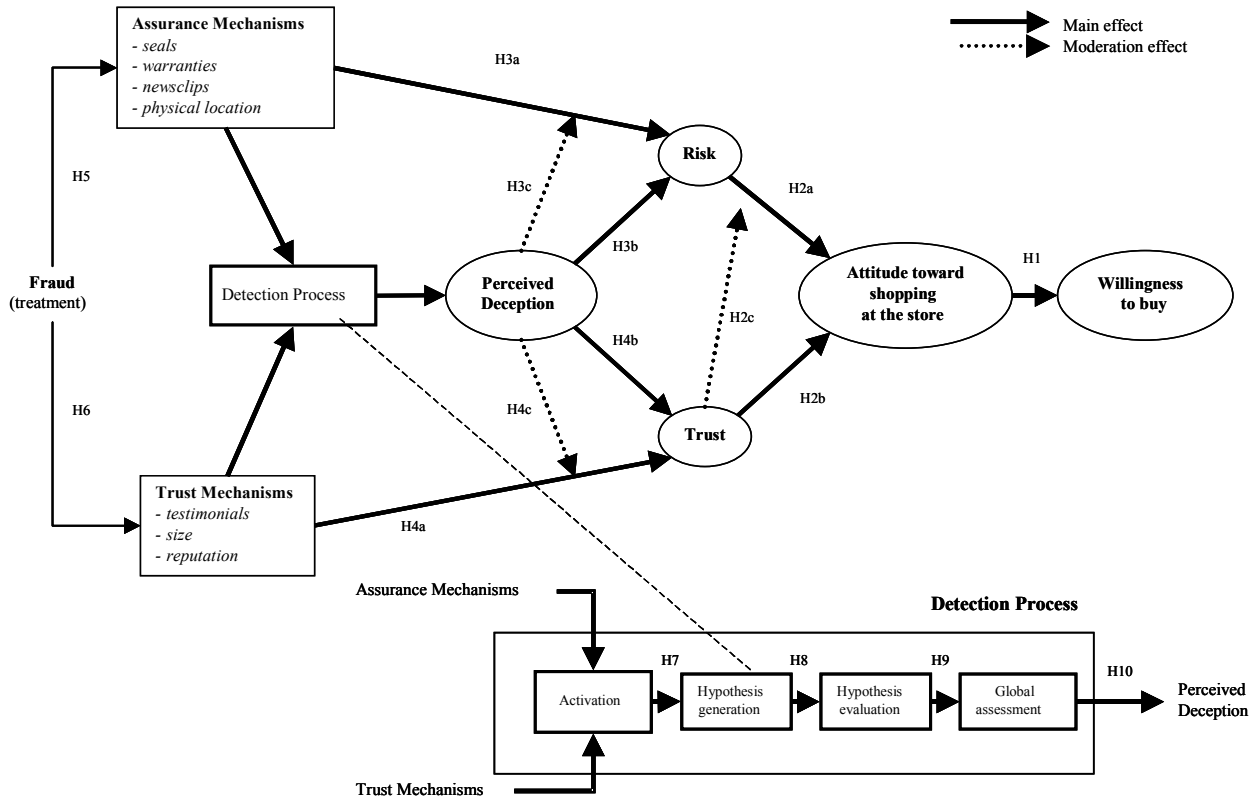


Figure 1. Extended DTR Model of Internet Consumer Behavior
(from Grazioli and Jarvenpaa 2000; Johnson et al. 2001)

Assurance mechanisms influence the perceived level of risk for a web store (hypothesis H3a). However, if a consumer suspects that the assurance mechanisms may have been deceptively manipulated, the influence of these mechanisms on perceived risk will be decreased (H3c). Furthermore, we expect that this perception of deception directly influences the perceived risk of the web store (H3b).

Trust-building mechanisms are means of fostering trust. DTR identifies three types of trust mechanisms: (1) web store reputation, (2) customer testimonials, and (3) web store size. Reputation is the extent to which buyers believe a web store is honest and concerned about its customers (Doney and Cannon 1997; Landon and Smith 1997). Reputation and customer testimonials convey information about the store’s past performance with other buyers. Large size suggests that other buyers trust the store and have successfully conducted business with it.

Trust mechanisms are expected to increase the level of consumer trust. We hypothesize that the level of trust is determined by reliance on trust mechanisms and by perceived deceptiveness (hypothesis H4a and H4b; see Figure 1). We also hypothesize that perceived deceptiveness moderates the relationship between trust mechanisms and risk (H4c). The suspicion that trust mechanisms may have been intentionally manipulated is likely to decrease the reliance on them.

Finally, we expect that malicious manipulations designed to introduce fictitious assurance and trust mechanisms or to increase the vividness of existing ones will have the effect of increasing the reliance on assurance (H5) and trust mechanisms (H6).

The Deception Detection Model

The theory of deception proposes an information-processing model of detecting deception. The model has been empirically validated in information-intensive domains (Johnson et al. 2001). The model is composed of four steps, presented in the lower right corner of Figure 1.

Activation. The potential victim of deception begins by identifying inconsistencies in the web store (e.g., “the link to the source of this newsclip is broken”). The deception detection model assumes that individuals continually compare the information cues they glean from a web store with their expectations about these cues. Both DTR and the theory of deception assume that individuals form expectations based on their previous (on- and off-line) experiences. When these expectations and the information from the web store do not match, an inconsistency is found.

According to the theory of deception, the identification of inconsistent cues (e.g., too-good-to-be-true warranties, seemingly exaggerated newsclips) is a requirement for success. If nothing anomalous is found, it is hard to conclude that a site is deceptive. A superior competence at looking for inconsistencies could, therefore, be a basis for superior fraud detection performance.

However, previous empirical results by Johnson and his colleagues suggest that successful and unsuccessful detectors do not differ in terms of ability to examine cues and identify inconsistencies. Anomalies are noticed by successful and unsuccessful detectors alike. For this reason, we do not expect that successful detectors inspect more cues than do unsuccessful detectors (H7).

Hypothesis generation. Sometimes, inconsistencies are ignored. More often, however, individuals offer interpretive hypotheses that explain the difference between expectations and observed cues (e.g., “the newsclips have been doctored,” “the seal is forged”). According to the theory of deception, these hypotheses are retrieved from templates stored in memory. Experiments with experienced professionals (auditors, loan officers) have suggested that hypothesis generation is crucial to successful detection. If a victim does not consider the possibility that a certain information item is deceptive, she cannot detect it. We, therefore, hypothesize that successful detectors are more likely to generate the hypothesis that an assurance mechanism or a trust mechanism has been manipulated (H8a and H8b).

Hypothesis evaluation. Once a hypothesis is generated, it needs to be evaluated. It might be necessary to choose among alternatives, or simply to decide whether to accept it or not. The theory of deception suggests that evaluation is done by comparison to some criterion. In domains where experience and feedback are abundant, it may be easier to develop valid criteria. Evaluating whether to accept a hypothesis might be problematic because a victim might not be capable of deciding whether “the newsclips have been doctored,” or “the seal is forged.” We, therefore, hypothesize that successful detectors are better able to successfully evaluate the hypothesis that assurance mechanisms and trust mechanisms have been manipulated (H9a and H9b).

Global assessment. The last step to reach the conclusion that a web store is deceptive consists of combining the available accepted hypotheses. The theory has emphasized that the assessment of deception can be the result of either a single strong hypothesis or the result of several weaker ones. We hypothesize that successful detectors use all available hypotheses (hypotheses about assurance mechanisms *and* about trust mechanisms) to generate an overall perception of the deceptiveness of a site (H10).

METHOD

A laboratory experiment tested the hypotheses identified above. This section describes the subject sample and the experimental design.

Subjects. A total of 93 undergraduate students participated in the experiment (35 males, 58 females). Subjects were volunteers recruited from several sections of a marketing class taught at a large U.S. university. As an incentive for their participation, the subjects received extra class credit.

All subjects were computer literate and familiar with Internet browsing; 82% reported browsing at least once a day, 88% had already bought at least one item on-line. The qualifier “unsophisticated” does not mean “beginner” or “novice” and was chosen to contrast the subjects of this study with the technology-savvy “technoMBAs” of the Grazioli and Jarvenpaa study.

Task. All subjects began the experiment by answering general questions about their attitudes toward on-line shopping. The main task consisted of asking all subjects to assume that a friend of theirs (“Jason Meeks”) has decided to buy a collectible doll for his girlfriend’s birthday. Browsing the web, Jason has identified a commercial web store that sells the collectible doll that his girlfriend wants (a \$350 limited-edition Barbie Doll). The doll has the right characteristics and price, but Jason is not completely comfortable about purchasing it from an on-line site. Therefore, he is asking the subject for a second opinion *on the site*.

The subjects accessed the web site from computers in a familiar computer lab at the university. They executed the task at their own pace. After examining the site, the subjects were asked to close their browser window and to fill in the rest of the questionnaire.

The scenario and task used in the experiment have several properties that make it attractive in terms of experimental realism: The site is a real commercial web store that agreed to collaborate in the study. The site is small and does not enjoy a recognizable brand name.

Design and Manipulations. Participants were randomly assigned to one of two conditions. Half of the subjects accessed the real commercial site, and the other half accessed a copy of the site forged by the experimenters. The forged site is an exact copy of the real site, except that it contains several malicious manipulations, designed to increase trust, decrease perceived risk, and ultimately increase the likelihood that visitors would evaluate the site positively.

Three manipulations affected assurance mechanisms: (1) A well-known third-party seal was accurately reproduced and prominently displayed on the site (Better Business Bureau—BBB On-Line). The seal was linked to a faithfully reproduced report by the BBB stating that the company's business record was clean. The forged URL of the report suggested that the report was stored on a BBB server. The report was also linked to the actual BBB site. (2) The warranty statement was modified, making it unrealistically liberal (e.g., "full refund," "no questions asked"). (3) Highly complimentary news clips from trade publications were forged. The quotes were linked to the web sites of the trade publications.

Manipulations 4 through 6 affected trust-building mechanisms: (4) A generic picture of a store was inserted on the site with a caption stating that the company had a physical location, which is not true. A fictitious Dallas address was given. (5) Site sales were grossly exaggerated (by two orders of magnitude). (6) Fictitious and enthusiastic endorsements from customers were also created and prominently displayed.

Detectability of Deception. Each of the manipulations introduced inconsistencies that can be detected and can arouse suspicion. (1) A moderate amount of exploration of the actual BBB site (reachable from the report mentioned above) leads to the BBB searchable database. Searching for the name of the company reveals that the company is not affiliated with the BBB. (2) The warranty is simply "too good to be true" according to common business practices. (3) Following the links to the quoted trade publications and searching for the quote proves unfruitful. (4) The provided picture of the Dallas store does not contain any recognizable sign or posting. The phone number has a California area code. (5) The e-mail addresses of two quoted customers turn out to be non-existent when searched for in a popular free e-mail service. Two other customer names appeared to be linked to personal web pages on popular free hosting services. One link is broken. The other connects to an unrelated page. (6) The whole inventory of the store includes a mere 15 dolls. However, based on the numbers prominently displayed on the site, the site should sell well over 100 collectible dolls a day, 365 days a year.

Priming manipulation. One third of the subjects in each condition received an abridged version of a Federal Trade Commission pamphlet on Internet fraud. Another third received the pamphlet and a watch list of possible manipulations. The last third did not receive any information. Subjects read the pamphlet before examining the site and answered 14 general questions designed to assess general comprehension (a manipulation check). The logic of this manipulation is described below.

Measures. The questionnaire filled by the subjects is closely based on the items and questions published in peer-reviewed work (Grazioli and Jarvenpaa 2000). **Attitude toward the store** was measured by asking the subjects whether it was likely that they would "return" to it, "tell another friend," or "consider" it for a purchase (three items). **Trust** was measured by asking whether the subject felt that the store was "trustworthy," "keeps its promises," "keeps customer best interest in mind," and "can be relied upon" (four items). **Perceived risk** was measured by asking whether the subjects felt that purchasing dolls at the site was "risky" (one item; Drolet and Morrison 2001). **Perceived deception** was measured by asking how "deceptive," "misleading," and "distorted" the information on the site was (three items). **Reliance on assurance mechanisms and on trust building mechanisms** was measured by asking subjects whether seals, news clips, warranties, and testimonials were "convincing," "believable" and "impartial" (four mechanisms x three items). **Perception of physical existence** was measured by asking whether the site "only existed on the web" and "has a physical presence" (two items). **Perceived size** was measured by asking whether the size is "large," is "one of the largest on the web," and whether it "is a small player" (three items).

Control variables included measures of computer self-efficacy (Compeau and Higgins 1995), attitude toward computers, toward web shopping, toward web safety, and toward risk.

Factor analyses were conducted to verify that the items included in each construct loaded as expected and without strong cross loading (Stevens 1996). The values of the constructs were computed as the mean of the seven-point ratings of the items associated with each construct. Table 1 contains descriptive statistics. No reliability was less than 0.78 (assurance mechanisms had an alpha of 0.85 in the cited study, which used identical items).

Table 1. Summary of Descriptive Statistics for the Main Constructs

Construct	Mean (1-7 range)	Standard Deviation	Reliability (Cronbach's alpha)
Willingness to buy from the web store	3.07	1.52	0.87
Trust	4.03	1.52	0.95
Perceived deception	3.36	1.45	0.93
Assurance-building mechanisms (seals, warranties, news clips, physical location)	4.62	1.06	0.88
Trust-building mechanisms (testimonials, size, reputation)	3.01	0.92	0.78
CONTROL VARIABLES			
Attitude toward the Web	3.53	1.24	0.72
Attitude to trust web stores	3.91	1.11	0.84
Positive attitude toward Internet safety	4.33	1.55	0.91
Positive attitude toward virtual stores	5.21	1.15	0.86

RESULTS

Hypothesis testing was conducted by linear and logistic regressions and by ANOVAs. Normality of residuals was tested and transformations applied where appropriate.

Detection success and failure. As measured by their behavior (recommending the purchase), most of the subjects (60.3%) failed. Of these, 30 missed the deception and felt comfortable with the fraudulent site while 26 issued false alarms (believed that the site was deceptive when it was not). Only 37 individuals were successful: 18 (19.4%) detected the manipulation and 19 (20.4%) correctly concluded that the site was clean. The distribution of successes and failures suggests that the subjects were not simply guessing (chi square $p = 0.04$), and that detection is possible, albeit difficult. These results are consistent with the previous study, as well as with predictions from deception theory.

Recommending the purchase and attitudes toward the store. DTR hypothesizes that the recommendation to buy depends on attitude toward the store (H1). A test based on a logistic regression accepts the hypothesis with $p < 0.000$ (see Table 2). The logistic model correctly predicts 83.9% of the observed behavior. Both findings are consistent with the results from the Grazioli and Jarvenpaa study. No control variable was found significant.

Attitude toward the store, trust and risk. Regression #2 tested whether positive attitude toward the store depend on trust and perceived risk, and whether trust moderates the relationship between risk and the attitude toward the store. The analysis concluded that trust increases positive attitudes toward shopping (H2a, $p = 0.002$). The standardized coefficient for trust is large and with the expected positive sign. Perceived risk affects negatively and significantly the attitudes toward the store (H2b, $p = 0.036$). No control variable was found significant.

A two-step procedure was employed to ascertain the moderating effects of trust on the relationship between risk and attitude toward the store. First, we created a dummy variable (HITRUST). HITRUST is equal to 1 if trust is above its mean (4.03), and 0 otherwise. The value of perceived risk was then multiplied to HITRUST. Ignoring the other variables in the model, the resulting regression equation is

$$\text{ATTITUDE} = \text{constant} + \beta_1 \text{RISK} + \beta_2 \text{RISK_BY_HITRUST} \quad (1)$$

Table 2. Regression Results

<p>Regression #1 (logistic) Hypothesis H1: Accepted Dependent Willingness to Recommend Purchase. Approx. R² = 0.64 Intercept *** Attitude toward shopping at store ***</p> <p>Regression #2 Hypothesis H2a, H2b: Accepted Hypothesis H2c: Rejected Dependent Attitude toward shopping at the store. Adj. R² = 0.49 Intercept n.s. Trust 0.482 ** Risk -0.208 * Risk by HITRUST n.s.</p> <p>Regression #3 Hypothesis H3b, H3c: Accepted H3a: Marginal (p=0.099) Dependent Perceived Risk Adj. R² = 0.53 Intercept ** Deception 0.315 * Assurance Mechanisms -0.175 marginal Assurance Mechanisms by HIDECEPTION 0.385 ***</p>	<p>Regression #4 Hypothesis H5a, H5b: Accepted Hypothesis H5b: Rejected Dependent Trust Adj. R² = 0.66 Intercept *** Trust Mechanisms 0.291 *** Deception -0.628 *** Trust Mechanisms by High Deception n.s.</p> <p>Regression #5 Hypothesis H6: Accepted Dependent Reliance on Trust Mechanisms Adj. R² = 0.11 Intercept *** Treatment (fraud) 0.343 ***</p> <p>Regression #6 Hypothesis H7: Accepted Dependent Reliance on Assurance Mechanisms Adj. R² = 0.14 Intercept *** Treatment (fraud) 0.385 ***</p>
--	--

NOTES:

1. Values represent standardized regression coefficients. Asterisks represent significance levels: ***p < 0.001; **p < 0.01; *p < 0.05; n.s. = non significant
2. Regression #1 is a logistic regression. All others are linear regressions. Logistic regression does not have an Adj. R² or standardized regression coefficients readily comparable to the linear regression coefficients.

or, equivalently

$$\text{ATTITUDE} = \text{constant} + \beta_1 \text{RISK} \quad \text{when trust is low, and} \quad (2)$$

$$\text{ATTITUDE} = \text{constant} + (\beta_1 + \beta_2)\text{RISK} \quad \text{when trust is high} \quad (3)$$

β_2 can be interpreted as the change in impact of perceived risk on attitude when trust is high. The results from regression #2 led to rejecting H2c (p = 0.862). When trust is high, individuals do not reduce the impact of risk considerations on the formation of their attitudes toward the web store. This contrasts with the results of the Grazioli and Jarvenpaa study. A plausible explanation is that it may take consumers some time to learn the relatively more sophisticated strategy of discounting risk in presence of trust. Only the more knowledgeable Internet consumers in the cited study had the opportunity to learn it.

Effect of the Assurance Mechanisms and Deception on Perceived Risk. Regression #3 tests the effects of the manipulated assurance mechanism on perceived risk, and used the same two-step procedure to test the direct and mediating effects of perceived deception on perceived risk (see Figure 1). The analysis confirmed that perception of deception significantly increases risk (H3b, p = 0.036), and has marginally confirmed that reliance on assurance mechanisms decreases risk (H3a, p = 0.099).

More interesting is the result on the moderating effect of deception (H3c, $p = 0.001$). When deception is high, there is no relationship between the assurance mechanisms and perceived risk. When consumers perceive that the site is deceptive, they strongly decrease their reliance on (possibly manipulated) assurance mechanisms.

Effect of the Trust Mechanisms and Deception on Trust. In a similar fashion, regression #4 tested the effects of the manipulated trust mechanisms on trust, and tested the direct and mediating effects of perceived deception. The analysis confirmed the expectations: trust is increased by trust mechanisms (H5a, $p < 0.000$) and decreased by deception (H5b, $p < 0.000$). As in the Grazioli and Jarvenpaa study, deception did not have a significant moderating effect on the relationship between trust mechanisms and trust (H5c, $p = 0.915$).

Effects of the deceptive manipulations. Regression analyses #5 and #6 show that the deceptive manipulations had the malicious effect of increasing reliance on assurance and trust mechanisms ($p = 0.001$ and $p = 0.000$, respectively).

Determinants of Success and Failure

This section compares the responses of those who either chose not to recommend purchasing from the fraudulent site (hits) or recommended purchasing from the clean site (correct rejections), with the responses of those who recommended purchasing from the fraudulent site (misses) or recommended against the clean site (false alarms).

Activation. To verify whether successful detectors examine more cues than non-successful detectors, we counted the number of the manipulated items that subjects reported having heeded (a series of yes/no questions). Since the fraudulent site contained more trust and assurance mechanisms than the original site, we compared only responses from subjects that examined the fraudulent site. As hypothesized, successful and unsuccessful detectors heeded the same number of items (H7, ANOVA $p = 0.608$; mean values 2.3 and 2.5, respectively). This is consistent with deception theory and suggests that competence at detecting deception does not lie in searching more intensely.

Hypothesis generation. To verify whether generation of the hypothesis of deception facilitates successful detection, we used a priming manipulation. Priming effects are well documented (e.g., Higgins et al. 1985). They are based on the idea that human memory contains cognitive structures that can be activated by certain stimuli (e.g., a printed word) and that when a structure is activated, it remains available for use in a subsequent task. For instance, you may be more likely to answer “FBI” than “IRS” to the question “name a Federal agency” because the FBI was mentioned at the beginning of this paper.

The logic of the test of H8 is as follows: we assume that the subjects who read the pamphlets about Internet fraud just before evaluating the web site were primed for the hypothesis that information on a web store can be deceptive. If generating the hypothesis that the web site is deceptive is a determinant of success at detecting deception, then these subjects should be more successful than subject who were not primed.

The test revealed no significant difference between the primed groups and the control group in terms of overall success, or in terms of perceived deceptiveness of either the assurance or trust building mechanisms. H8a and H8b were both rejected. This result contradicts studies conducted with experienced professionals (e.g., Johnson et al. 1992). A possible explanation of the divergence is that the experimental stimulus was in some way lacking. However, we used official material developed for the specific purpose of facilitating fraud detection by consumers. We also ran a comprehensive and successful manipulation check. Below we further discuss this result.

Hypothesis evaluation. To verify whether the evaluation of the hypothesis of deception facilitates successful detection, we compared the ratings given by successful and unsuccessful subjects to the assurance and trust mechanisms.

Hypothesis H9a and H9b were both accepted. ANOVA of the data concluded that successful subjects rated their reliance on the trust and assurance mechanisms significantly lower in the fraudulent site ($p < 0.000$ and $p < 0.000$) than unsuccessful subjects. They also rated the reliance of the trust mechanisms significantly higher in the clean site ($p = 0.029$). No significant difference was found with respect to the assurance mechanisms for the clean case ($p = 0.101$).

The last two series of tests can be interpreted by stating that the generation of the hypotheses that assurance or trust mechanisms might be manipulated is not a determinant of success for unsophisticated consumers because they do not have the capacity to effectively assess these hypotheses. Those who can evaluate assurance or trust mechanisms succeed at the overall task,

independent of priming. In other words, priming subjects to facilitate hypothesis generation is useless, if we do not also provide them with effective means to assess the primed hypothesis.

Global assessment. To verify how successful detectors use all available hypotheses about assurance and trust mechanisms to form an overall perception of deceptiveness of a site (H10), we regressed perceived deception on reliance on the assurance and on the trust mechanisms. We did this separately for the successful and unsuccessful individuals.

The results do not support H10, in an interesting way. The successful individuals include the assurance mechanisms ($p = 0.005$), but strongly discount the trust mechanisms ($p = 0.419$) as they assess the deceptiveness of the site. The *unsuccessful* individuals include both assurance mechanisms ($p = 0.000$) and trust mechanisms ($p = 0.002$). It seems that discounting the trust mechanisms, which are perhaps perceived as easier to manipulate, is a characteristic of successful detectors.

Limitations

Perhaps the chief limitation of this study is that the subjects did not actually purchase from the store: they advised a hypothetical friend about purchasing. It might be that individuals with more at stake (i.e., their own money) would work harder at evaluating the site and perhaps be more successful. The amount of time spent inspecting the site and the enthusiasm demonstrated by some of the subjects suggests, however, that this might not be a major threat to validity.

A second limitation is that our results apply to a medium-price specialty item and might not generalize to other consumer goods. A mitigating consideration is that our results are generally consistent with the results of the Grazioli and Jarvenpaa experiment, where the item of the purchase was a medium-high price item (a laptop).

A third limitation is that our results are likely to generalize to small and relatively unknown web stores, not to well-established brand names. This is not because brand names cannot be subject to page jacking, but because the effect of familiarity with the site might override the perception of possible deception.

A fourth limitation is that we have examined responses to one particular type of fraud (page jacking), which is arguably hard to detect because it is relatively new. Detecting more familiar types of deception (e.g., a pyramid scheme) might be easier because consumers might more readily generate hypotheses about it.

A last limitation is that we focused on deliberate, problem-solving aspects of detecting deception and ignored the affect components that are likely to be involved in the decision to buy. In addition, our results might be culturally biased, to the extent that individuals in the United States have a disposition to trust strangers and try the new.

CONCLUSIONS

This work contributes to the ongoing scientific debate on trust, deception, and risk on the Internet by integrating two streams of empirical research. We have proposed an extension to the DTR model and developed several new hypotheses from it. A lab experiment has successfully replicated past results with new subjects and obtained new findings.

Previous laboratory work (Grazioli and Jarvenpaa 2000) looked at seasoned Internet consumers, individuals whose knowledge of e-commerce technology and business made generalizations to the larger population of Internet users problematic. The cited study warned that if those Internet consumers can be deceived, the risk for the larger public might be even greater. This work has validated that suspicion by employing a sample that better approximates the wider population of Internet users.

We added to the previous work a new focus on the process of detecting deception. The picture of the Internet consumer that emerges from this new focus suggests that many unsophisticated users are struggling to develop effective strategies to transact in a virtual world. When compared to more successful and more mature consumers, they insufficiently discount risk in the presence of trust, are unable to evaluate trust building and assurance mechanisms that they identify on a web site, and are unable to effectively combine the information they gather.

The results of this work suggest a number of worthwhile directions for future research: studies more explicitly designed for capturing process traces (e.g., verbal protocols or click-through analyses) could confirm and sharpen our conclusions about

detection success and failure. There might be different detection strategies. One reviewer of this work suggested that trust-building and assurance mechanisms might interact in influencing the outcomes of the detection process.

At the beginning of this paper, we stated that this research has been motivated by the concern expressed by several government agencies over Internet fraud. We have found that consumers are in fact at risk. At the same time, it seems that progress is possible if we find ways to educate individuals to better evaluate the clues to deception. They already look at them. They just do not see them for what they are.

References

- Compeau, D., and Higgins, C. "Computer Self Efficacy: Development of a Measure and Initial Test," *MIS Quarterly*, Fall 1995, pp. 198-211.
- Crosby, L. A., Evans, K. R., and Cowles, D. "Relationship Quality in Services Selling: An Interpersonal Influence Perspective," *Journal of Marketing* (53), July 1990, pp. 68-81.
- Das, T. K., and Teng, B. S. "Between Trust and Control: Developing Confidence in Partner Cooperation in Alliances," *Academy of Management Review* (33:3), 1998, pp. 491-512.
- Dasgupta, P. "Trust as a Commodity," in *Trust: Making and Breaking Cooperative Relations*, D. Gambetta (ed.), Basil Blackwell, New York, 1998, pp. 47-72.
- Doney, P. M., and Cannon, J. P. "An Examination of the Nature of Trust in Buyer-Seller Relationships," *Journal of Marketing* (61), April 1997, pp. 35-51.
- Dowling, G. R., and Staelin, R. "A Model of Perceived Risk and Intended Risk-Handling Activity," *Journal of Consumer Research* (21), June 1994, pp. 119-134.
- Drolet, A., and Morrison, D. "Do We Really Need Multiple Items Measures in Service Research," *Journal of Service Research* (3:3), 2001, pp. 196-204.
- Ekman, P. *Telling Lies: Clues to Deceit in the Marketplace, Politics, and Marriage*, W. W. Norton & Company, New York, 1992.
- Federal Trade Commission, "Prepared Statement of the FTC on Fraudulent, Marketing Schemas before the Subcommittee on Commerce, Justice, State, and the Judiciary of the Senate Appropriations Committee," United States Senate, 5 February 1998.
- Federal Trade Commission v. Pereira et al. "Complaint for Permanent Injunction and other Equitable Relief," Case No. 99-1367-A, U.S. District Court, E. D., Alexandria, VA, 1999.
- Fishbein, M., and Ajzen, I. *Belief, Attitude, Intention and Behavior: An Introduction to Theory and Research*, Addison-Wesley, Reading, MA, 1975.
- Ganesan, S. "Determinants of Long-Term Orientation in Buyer-Seller Relationships," *Journal of Marketing* (58:2), April 1994, pp. 1-19.
- Grazioli, S., and Jarvenpaa S. "Perils of Internet Fraud," *IEEE Transactions on Systems, Man, and Cybernetics* (30:4), 2000, pp. 395-410.
- Higgins E. T., Bargh, J. A., and Lombardi, W. "The Nature of Priming Effects in Categorization," *Journal of Experimental Psychology: Learning, Memory and Cognition* (11), 1985, pp. 59-69.
- Jarvenpaa, S. L., and Grazioli, S. "Surfing Among Sharks: How to Gain Trust in Cyberspace," *Financial Times*, Mastering Information Technology Section, March 15, 1999, pp. 2-3.
- Johnson, P. E., Grazioli, S., and Jamal, K. "Fraud Detection: Intentionality and Deception in Cognition," *Accounting, Organization and Society* (18:5), 1993, pp. 467-488.
- Johnson, P. E., Grazioli, S., Jamal, K., and Berryman, G. "Detecting Deception: Adversarial Problem Solving in a Low Base Rate World," *Cognitive Science* (25:3), 2001, pp. 355-392.
- Johnson, P. E., Grazioli, S., Jamal, K., and Zualkernan, I. A. "Success and Failure in Expert Reasoning," *Organizational Behavior and Human Decision Processes* (53:2), 1992, pp. 173-203.
- Kelley, H. H., and Thibaut, J. W. *Interpersonal Relations: A Theory of Interdependence*, Wiley, New York, 1978.
- Kollock, P. "The Emergence of Exchange Structures: An Experimental Study of Uncertainty, Commitment, and Trust," *American Journal of Sociology* (100), 1994, pp. 313-345.
- Kramer, R. M. "Trust and Distrust in Organizations: Emerging Perspectives, Enduring Questions," *Annual Review of Psychology* (50), 1999, pp. 569-598.
- Landon, S., and Smith, C. E. "The Use of Quality and Reputation Indicators by Consumers: The Case of Bordeaux Wine," *Journal of Consumer Policy* (20), 1997, pp. 289-323.

- Macintosh, G., and Lockshin, L. W. "Retail Relationships and Store Loyalty: A Multi-level Perspective," *International Journal of Research in Marketing* (14:5), December 1997, pp. 487-497.
- McKnight, D. H., Cummings, L. L., and Chervany, N. L. "Initial Trust Formation in New Organizational Relationships," *Academy of Management Review* (23:3), 1998, pp. 473-490.
- Mitchell, R. W., and Thompson, N. S. (Eds.). *Deception: Perspectives on Human and Non-human Deceit*, SUNY Press, Albany, NY, 1986.
- Morris-Cotterill, N. "Use and Abuse of the Internet in Fraud and Money Laundering," *International Review of Law Computers and Technology* (13:2), 1999, pp. 211-228,.
- Pitofsky, R. "Prepared Statement of the Federal Trade Commission on 'Internet Fraud' Before the Subcommittee on Investigations of the Governmental Affairs Committee United States Senate," Washington, DC, February 10, 1998.
- Spreitzer, G. M., and Mishra, A. K. "Giving Up Control without Losing Control," *Group and Organization Management* (24:2), 1999, pp. 155-187.
- Stevens, J. *Applied Multivariate Statistics for the Social Sciences*, Lawrence Erlbaum Associates, Inc., Mahwah, NJ, 1996.
- Thagard, P. "Adversarial Problem Solving: Modeling an Opponent Using Explanatory Coherence," *Cognitive Science* (16), 1992, pp. 123-149.

