working draft

Adverse Selection in Online "Trust" Certifications

Benjamin Edelman
Harvard University
edelman@pobox.com

May 11, 2006

Abstract

Widely-used online "trust" authorities issue certifications without substantial verification of the actual trustworthiness of recipients. Their lax approach gives rise to adverse selection: The sites that seek and obtain trust certifications are actually significantly less trustworthy than those that forego certification. I demonstrate this adverse selection empirically via a new dataset on web site characteristics and safety. I find that TRUSTe-certified sites are more than twice as likely to be untrustworthy as uncertified sites, a difference which remains statistically and economically significant when restricted to "complex" commercial sites. I also present analogous results of adverse selection in search engine advertising – finding ads at leading search engines to be more than twice as likely to be untrustworthy as corresponding organic search results for the same search terms. However, I find no sign of adverse selection in search engines' organic results; on the whole, the safety of organic results matches the safety of the web at large.

Keywords: Adverse selection, certification, reputation, trust, Internet, search engines.

I thank seminar participants at Harvard University's Department of Economics, Business School, and Department of Computer Science. I am grateful to Robert Akerlof, Peter Coles, Chris Dixon, David Parkes, and Stuart Schechter for helpful comments and advice.

1 Introduction

When agents have hidden types, contract theory warns of bad results and potentially even market unraveling. Since Akerlof's defining "lemons" (1970), others have worried about similar problems in other markets – such as bad drivers wanting more car insurance than good drivers (Chiappori and Salanie 2000), and healthy people disproportionately buying annuities (Finkelstein et al, 2004).

In general, it is difficult to empirically assess the significance of adverse selection problems. For example, used car markets are made more complicated by idiosyncratic details – unobservable car characteristics, local markets, and casual sellers. Some work manages to address these problems. For example, Chiappori and Salanie focus on novice drivers, who have less private information about their own type (since they have not yet started to drive), such that economists can observe most relevant characteristics. But these special cases bring problems of their own. Researchers may be less interested in the absence of adverse selection among novice drivers' insurance purchases, and more interested in the adverse selection that (perhaps) actually does affect most other drivers.

This paper applies an adverse selection model to a new market – Internet web sites and their associated "trust"-type certifications. With a new data source, I observe and analyze variables generally unobservable both to consumers and to certification authorities. In the face of sites' hidden types, these certifications provide an appropriate area to test and measure theories of adverse selection.

Beyond adverse selection, trust certifications are also of interest in their own right. Such certifications have played an important role in the policy debate as to regulation of online privacy and safety, and typical Internet users see such certifications remarkably frequently. Yet, empirically, adverse selection turns out to be important in this market: My analysis indicates that adverse selection substantially reduces overall certification quality. In particular, I find that sites certified by one well-known authority, TRUSTe, are more than twice as likely to be untrustworthy as uncertified sites.

1.1. The Basic Web Site Safety Problem

Consumers seeking online information and services face a serious problem in deciding what firms to do business with. Users could stick with a few "known-good" household names, but such a narrow focus would reduce match quality, denying users access to much of the rich diversity of Internet content and services. Yet venturing into the unknown Internet carries important risks too: Untrustworthy sites might send users spam (if users register or otherwise provide their email addresses), infect users' computers with viruses or other harmful code (if users install the programs they offer), or simply fail to deliver the promised merchandise (if users make purchases). Ex ante, users have no easy way to know which sites to trust. A seemingly-innocent site could turn out to be a wolf in sheep's clothing.

These online interactions reflect a two-sided market – with sites actively making decisions about how to present themselves. Good sites want to demonstrate that they're good. But as in the usual moral hazard framework, bad sites pretend they're good too.

Facing numerous untrustworthy or even malicious sites, some analysts call for government regulation. In principle, a government agency might diligently review web sites in search of spam, scams, and harmful programs. To some extent, the FTC and state attorneys general perform such investigations – though it seems that their efforts address only a small portion of bad actors. As a practical matter, government intervention seems inept. See Tang et al. (2005), presenting a model of enforcement of online privacy breaches, finding mandatory government standards appropriate only for serious harms. So government intervention does not appear to be a realistic way forward, in general.

At the other extreme, users might be left entirely on their own. In an extreme laissez faire world, no one vouches for anything found on the Internet, and no one cleans up the inevitable messes: Users won't get help from regulators or computer makers, and perhaps not even from their own IT departments. So laissez faire leaves users with whatever personal defenses they can muster and with whatever protective technologies they can purchase or devise. In some respects, laissez faire is a reasonable description of the current state of affairs. (IT departments can't protect users from getting ripped off. And many an IT department feels powerless to stop spam.) But the laissez faire approach carries substantial costs – leading users to take excessive precautions, preventing the formation of otherwise-profitable relationships. Users would buy more products, join more sites, and download more programs were it not for their well-founded fears in a laissez faire regime.

Finally, there exists a middle approach between the extremes of government regulation and laissez faire: A non-governmental rating organization. Such an organization would identify specific bad behaviors, then evaluate sites' compliance with those baseline standards. If compliance evaluations were accurate and low-cost, this authority might support an equilibrium where all good firms receive positive ratings, and where all consumers use only those sites that receive positive ratings. Tang et al. describe this approach as appropriate for a broad class of online behaviors. But even so, there are reasons to worry of its effectiveness. Extending sites to take a continuum of types, a single binary certification may not convey all necessary information about all possible site behaviors. Insufficient precision is particularly likely if consumers are heterogeneous in preferences, especially in their assessments of objectionable behaviors. Finally, it is hard to identify what specific behaviors are "bad," particularly when firms have strong economic incentives to blur the boundaries. So in practice, a non-governmental rating authority may be less expedient than it initially seems – including for the reasons set out in subsequent sections.

1.2. Certification Authorities

Most prominent among non-governmental rating organizations are so-called "trust" certification authorities. These rating organizations set out specific criteria for membership, often focusing on privacy or on safety more generally. The organizations

reward their members by offering seals to be placed on recipients' web sites, often at the point where site operators want to reassure users of sites' legitimacy and trustworthiness (e.g. at registration forms and checkout pages). Particularly well-known certification authorities are TRUSTe and BBBonline.

In principle, certification authorities might set and enforce substantive and procedural provisions sufficiently rigorous that certified members are highly likely to satisfy reasonable consumers' expectations of safety. But in practice, there is reasonable basis to doubt the effectiveness of certain certification authorities. LaRose and Rifon (2002) offer a stinging critique: Certification authorities have granted multiple certifications to firms under investigation by the FTC for privacy policy violations; certification authorities have failed to pursue complaints against major companies whose privacy breaches were found to be "inadvertent"; and in one case a certification authority even failed to abide by its own privacy policy. Singel (2006) also questions the effectiveness of TRUSTe's effort: In a 2004 investigation after user complaints, TRUSTe gave Datran Media a clean bill of health; yet a subsequent New York Attorney General statement and settlement indicated that Datran had committed the most far-reaching privacy policy violations known to date.

As a threshold matter, some certification authorities' substantive standards often seem substantially duplicative with companies' existing legal or practical requirements. Consider the requirements summarized in TRUSTe's "Program Requirements" document. For example, the first listed requirement, offering an email unsubscribe function, duplicates Sec.5.(a)(4)(A) of the federal CAN-SPAM Act. Similarly, the first security requirement, using SSL encryption or similar technology to protect sensitive information like credit card numbers, is already commercially widespread due to rules established by credit card networks. Other certification authority requirements may be somewhat more rigorous, but in general there is reason to doubt their significance. See also Boutin (2002), reporting TRUSTe initially lacking any *substantive* requirements whatsoever (i.e. requiring only the *presence* of a privacy policy).[1]

Tellingly, strikingly few certificates have been revoked. For example, TRUSTe's Fact Sheet (2006) reports only two certificates revoked in TRUSTe's ten-year history. Of course there are multiple explanations for an absence of enforcement actions. Suppose certificate recipients think a certification authority will detect infractions with a high probability and will extract a harsh punishment. Then that authority will attract only "good" sites, and the authority might never actually detect any wrongdoing, nor ever actually have to punish any recipient. But this tough-dog theory stands contrary to observed facts: For example, TRUSTe has only a small staff, with little obvious ability to detect violations of its rules. TRUSTe's posted rules and procedures reveal substantial focus on sites' self-certifications and on user complaints. Rule violations at TRUSTe member sites have repeatedly been found by independent third parties, not by TRUSTe itself.

[1] Some industry sources report TRUSTe claims that most applicants must make some change to their sites or their business practices in order to obtain TRUSTe certification. However, TRUSTe and its members do not report these changes, even in general terms. It is therefore difficult to assess their significance or to determine whether they provide any substantial benefits to users.

TRUSTe's "Watchdog Reports" page also indicates a lack of focus on enforcement. According to TRUSTe's posted data, users continue to submit literally hundreds of complaints each month. But of the 3,416 complaints received since January 2003, TRUSTe concluded that *not a single one* required any change to any TRUSTe member's operations, privacy statement, or privacy practices, nor did any complaint require any revocation or on-site audit. Other aspects of TRUSTe's watchdog system also indicate a lack of diligence.[2]

Finally, as Greenstadt and Smith (2005) point out, certification authorities are "captured" – paid by the same companies they certify. Certification authorities have little incentive to antagonize their customers: Any such pressure would harm the authority's profits by discouraging renewals and future applications.

Even the creators of certification authorities seem unhappy with their development. Boutin (2002) quotes TRUSTe co-founder Esther Dyson conceding that TRUSTe is "a little too corporate" and that TRUSTe lacks the "moral courage" to criticize violations. Similarly, TRUSTe co-founder Electronic Frontier Foundation admitted in a 1999 letter to the FTC that "legislation is needed" to protect users' privacy and that "it is time to move away from a strict self-regulation approach."

Facing allegations of low substantive standards, lax enforcement, and ethical compromise, it is unclear whether site certifications offer direct benefits to consumers. Furthermore, consumers are unlikely to place substantial value on certification authorities' offer to assist with dispute resolution: TRUSTe's lenient disposition of watchdog reports indicates that complaining consumers rarely cause changes in members' business practices.

Despite certification authorities' limited substantive role in assuring good online practices, at least some consumers seem to regard a certification authority's certification as a significant positive signal. For example, in promoting its service to potential applicants, TRUSTe touts its benefits to certificate recipient Realty Tracker, which says TRUSTe "convey[ed] trust" and "built confidence" with Realty Tracker's visitors, yielding "an increase in registrations." See LaRose and Rifon, characterizing certification authorities' seals as "persuasion attempt[s]."

Other firms are well-equipped to evaluate claims of benefits to certification: Firms could randomize their inclusion of TRUSTe or similar seals, thereby determining whether seals actually increase registrations and sales. In the related context of comparison shopping engines, Baye and Morgan (2003) empirically confirm the benefits of certification seals: Merchants with seals enjoy a price premium over uncertified merchants.

---

[2] Inexplicably, TRUSTe's last posted watchdog data report dates from mid-2004; more recent data is not available on TRUSTe's Watchdog Reports page, https://www.truste.org/consumers/watchdog_reports.php. Furthermore, the posted reports show a striking inattention to detail. For example, the first seven months' reports all bear the title "Watchdog Report for October 2000" (correct only for the first of those reports), and the next 37 reports all say "Watchdog Report for May 2001" (a title also inaccurate for all but the first).

Whatever the actual merits of certification authorities as arbiters of trust, some government authorities seem to regard certification authorities as an appropriate step forward. See the FTC's 1999 "Self-Regulation and Privacy Online," finding private-sector certification authorities preferable to direct FTC regulation.

The FTC specifically cites two well-known certification systems: TRUSTe's Web Privacy Seal and BBBOnLine's Privacy Seal Program. These certification authorities are the focus of my subsequent analysis, due to the prevalence, their relatively large member lists (compared to other certification authorities), and their decisions to post their member lists (providing data necessary for my analysis). I largely focus on TRUSTe, the first online trust certification authority, the largest, and (it seems) still the best-known.

1.3. Search Engines as Arbiters of Trust

Though less explicitly focused on trust, search engines also play a prominent role in influencing users' decisions as to what sites and services are safe to use. High placement at a search engine might reasonably be interpreted as a kind of endorsement: A search engine's high ranking indicates that a given site is (believed to be) among the best resources for a given search term. (Gaudeul 2004)

Empirical work demonstrates that users place substantial weight on high search engine rankings. Consumers believe highest-ranked sites are most likely to serve their interests (Marable 2003), and top-ranked sites have the highest click-through rates. (Joachims 2005) Because users may not understand the difference between paid search engine advertising versus ordinary "organic" listings (Consumer Reports WebWatch 2002), Marable's result is likely to apply to all search engine results, not just to organic results.

Search engines present two distinct notions of certification. First, a site might be said to be certified if it appears high in organic search results, rather than appearing further down or (perhaps) not at all. Second, a site might be said to be certified if it appears as a sponsored result, e.g. a search engine advertisement. Subsequent analysis will compare search engine trustworthiness along these two different axes of site placement.

In subsequent analysis, I largely group search engines together with certification authorities. I use the term "certification authority" to refer specifically to certificate-granting organizations such as TRUSTe, while I use the broader term "trust authority" to include search engines also.

2 Theory of Adverse Selection in Trust Authorities

Suppose that, as described above, certain trust authorities issue certifications of trustworthiness without rigorous assessment of the true trustworthiness of recipients. This market structure creates significant reason to worry of adverse selection in certification recipients. Certifications of trustworthiness seek to signal consumers that the certified firms are in fact highly likely to be trustworthy. But if untrustworthy firms

can obtain certifications just as easily as trustworthy firms, then consumers have little reason to conclude that a certified firm is trustworthy: Rational consumers would rightly worry that certified firms obtained certifications despite actually being untrustworthy.

To provide consumers with the intended signal, a trust authority's certification must increase a rational consumer's assessed probability that a given site is trustworthy. Suppose a rational consumer has some prior belief P(t) that a given site is trustworthy, before receiving a signal (denoted "**s**") of trustworthiness ("**t**"). Such a consumer should update his probability according to the usual Bayes Rule formula:

$$P(t|s) = \frac{P(s|t)\,P(t)}{P(s)} \tag{1}$$

Expanding the denominator using the Law of Total Probability:

$$P(t|s) = \frac{P(s|t)\,P(t)}{P(s|t)\,P(t) + P(s|\bar{t})\,P(\bar{t})} \tag{2}$$

For consumer's assessment of site trustworthiness to increase as a result of a site's certification, it must be the case that:

$$P(t|s) > P(t) \tag{3}$$

which implies:

$$\frac{P(s|t)}{P(s|t)\,P(t) + P(s|\bar{t})\,P(\bar{t})} > 1 \tag{4}$$

Rearranging further, using the fact that $P(t) = 1 - P(\bar{t})$:

$$P(s|t) > P(s|t)\,P(t) + P(s|\bar{t})\,P(\bar{t}) \tag{5}$$

$$P(s|t)\,(1-P(t)) > P(s|\bar{t})\,P(\bar{t}) \tag{6}$$

$$P(s|t)\,P(\bar{t}) > P(s|\bar{t})\,P(\bar{t}) \tag{7}$$

$$P(s|t) > P(s|\bar{t}) \tag{8}$$

Equation 8 offers an intuitive result: For a certification to cause a consumer to conclude a certified site is more safe than the consumer thought ex ante, the certification must be given to trustworthy sites more often than it is given to untrustworthy sites. Equation 8 also shows two other possible cases: If trustworthy and untrustworthy sites receive the certification equally frequently, i.e. $P(s|t)=P(s|\bar{t})$, then the certification is uninformative. Finally, if the certification is given to untrustworthy sites more than to trustworthy sites, i.e. $P(s|t)<P(s|\bar{t})$, then consumers would rationally consider a site *less* trustworthy when it reveals that it has been certified.

Equation 8 yields an empirical strategy for testing site certifications: Compare the certification rates of trustworthy sites with the certification rates of untrustworthy sites. Alternatively, further rearranging confirms that it is equivalent to compare the trustworthiness rates of certified sites, relative to the trustworthiness rates of uncertified sites. (See Appendix for proof.) For a valid certification that increases consumers' ex post assessment of site trustworthiness, certified sites must be more likely to be trustworthy than are uncertified sites. Formally, a valid certification requires

$$P(t|s) > P(t|\bar{s}) \tag{9}$$

The preceding adverse selection model offers a clear empirical prediction: That the inequality in (9) should fail. In particular, if adverse selection substantially affects these certifications, then certified sites should be *less* safe then certified sites.

Analyzing correlations between trustworthiness and certification is analogous to the approach in the existing adverse selection literature. Consider Finkelstein (2004), finding that annuitants live longer than non-annuitants. Intuitively, adverse selection implies that annuities are more appealing to individuals who know they will live a long time. Finkelstein shows exactly this result, finding a positive relationship between claimed type (annuity purchase) and outcome (lifetime). Chiappori and Salanie (2000) use the same method to demonstrate the absence of adverse selection in car insurance for novice drivers in France – finding no correlation between the conditional distributions of claimed type (insurance purchase) and outcome (insurance claims). Similarly, Genesove (1993) extends these correlations with the equilibrium assumption that average price in a given market must reflect average quality in that market. He then regresses auction bids on variables including a type-determining variable (there, whether a given used car was sold by a dealer who exclusively sells used cars, or by a dealer who largely sells new cars), interpreting a significant coefficient as evidence of adverse selection at used car dealers.

In the context of online trustworthiness, adverse selection offers a second empirical prediction: That adverse selection should worsen over time. When a new certification authority starts operation, its certificate of trustworthiness has no clear value. The certificate – ultimately merely a graphical image to be placed on recipients' sites – can easily be replicated by any other self-styled trust authority or rogue third party. Such a certification is of uncertain initial value to initial recipients. "Good-type" recipients may want the certification for some intrinsic reason, i.e. as a way to assist the formation of a trust authority they consider an important project. But "bad-type" recipients will have no reason to want the certification: Consumers initially will not know what the certification (purportedly) means, and consumers therefore will not defer to it. So in initial periods, only trustworthy firms should get certified. Later, once the certification gains traction, untrustworthy firms should join too. This theory yields a second hypothesis: That certification authorities should not suffer adverse selection in initial periods of operation, but that adverse selection should worsen over time.

My suggestion of adverse selection at online certification authorities is not without precedent. See LaRose and Rifon, finding that certified sites privacy policies allow more

invasive data collection than corresponding policies at uncertified sites. But where LaRose and Rifon draw conclusions based on a hand-scored assessment of 200 sites, I extend analysis to many thousands of sites using automated methods. I also consider axes of trustworthiness other than invasiveness of data collection permitted under sites' privacy policies.

2.1. Trust Authorities in Equilibrium

Critics of online trust authorities might reasonably wonder how trust authorities can exist in equilibrium. Suppose, as hypothesized above, that trust authorities suffer adverse selection – such that certified sites are actually less deserving of trust, on average, then uncertified sites. Alternatively, suppose trust authorities award certifications randomly, uncorrelated with sites' actual trustworthiness. In equilibrium users should learn that so-called trust certifications are uninformative, and users should therefore discount – ignore! – those certifications. But if consumers ignore the certifications, sites have no incentive to become certified. Then certification schemes should disappear altogether.

In principle, it is reassuring to see a prediction that worthless trust authorities will self-destruct. After all, a worthless signal is not worth sending. If firms decline to send worthless signals, then less economic resources will be devoted to any such waste.

But under reasonable assumptions, trust authorities might continue to exist in the long run. Suppose the Internet receives a continuous stream of naïve new users, who mistakenly think a site is truly trustworthy when it (or its certificate authority or a search engine) says so. If the number of such users is sufficient, sites might find it profitable to continue to present trust certificates, even though most users know the certificates are worthless. Especially if certifications cost little to obtain, the certificates will offer a measurable benefit with no substantial downside.

The prior section's second hypothesis offers another reason why certification authorities might survive. On that theory, trust authorities started out "good," with members that were truly trustworthy. This good period creates goodwill among sophisticated, informed consumers who, during this initial period, observe that a given trust authority has members who are actually trustworthy. That goodwill might take some time to dissipate, i.e. in the face of slow learning. So even if all consumers are sophisticated, slow learning (by at least some consumers) could also explain the survival of certification authorities.

Finally, exogenous market forces might preserve certain trust authorities. The large companies that founded TRUSTe are likely to continue to support it, so long as it serves their regulatory goals. So even if TRUSTe would otherwise face extinction, these core members may keep it afloat. Similarly, search engines are bound to issue implicit certifications, even if their certifications carry little weight. It seems search engines inevitably rank some sites above others – necessarily lending credibility to the sites placed in top positions. So trust authorities may also continue to exist for an extended period, even if basic economic theory suggests that they should disappear.

But even if trust authorities continue to exist for reasons exogenous to their signaling value, users might come to discount or ignore the authorities' certifications. Consider the plight of home shopping networks and late-night TV sales pitches: Home-shopping TV may still exist, but many consumers view these channels with disdain. Discredited trust authorities could face a similar plight.

2.2. Which Sites Seek Certification

The prior section suggests that certification authorities might continue in equilibrium because naïve users mistakenly believe authorities' certifications. But on that view, *all* sites should end up certified – contrary to empirical observation that many sites are uncertified. This section attempts to address that prediction via a simple algebraic model of site certification.

Suppose proportion $\mathbf{p_i}$ of users are naïve at a given site **i**, while the remaining $\mathbf{1-p_i}$ are sophisticated. Sophisticated users make purchases from that site with probability $\mathbf{q_i}$, and sophisticated users always ignore sites' certifications. But if a site is certified, naïve users purchase with a higher probability $\mathbf{q_i+r_i}$. (For naïve users, $\mathbf{r_i}$ thus gives the marginal increase in purchase probability when a site obtains certification.) A certificate costs **c** (per year). A given site receives $\mathbf{n_i}$ visitors per year and earns profit p_i from each sale. A rational site obtains certification if the certification increases profits, i.e.:

$$[\text{profits from certification}] > [\text{profits if not certified}] \tag{10}$$

Specifically:

$$\begin{aligned}&[\text{profits from naïve users if certified}]\\ &\quad + [\text{profits from sophisticated users if certified}]\\ &\quad - [\text{costs of certification}]\\ &\qquad\qquad > [\text{profits from all users if not certified}]\end{aligned} \tag{11}$$

Substituting:

$$[n_i\pi_i][p_i(q_i+r_i) + (1-p_i)q_i] - c > n_i\pi_i q_i \tag{12}$$

Rearranging and canceling, a site obtains a certification if:

$$n_i\pi_i p_i r_i > c \tag{13}$$

In general, the prior section assumed $\mathbf{n_i}$, p_i, $\mathbf{p_i}$, and $\mathbf{r_i}$ were all strictly positive, and that **c** was relatively small. These are good assumptions for typical commercial sites, with substantial naïve users (**p** nonnegligible) who are substantially confused about the value of certifications (**r** nonnegligible), and with certification low cost relative to site size (**c** small). The model predicts that all such sites should seek and obtain certifications.

But consider a site with more sophisticated users. Their sophistication could enter via a sufficiently small $\mathbf{p_i}$ – a site where very few users are naïve, i.e. because they understand the true meaning of a site's certification. Alternatively, their sophistication

could enter via a smaller $\mathbf{r_i}$ – no benefit to naïve users' confusion about certification, i.e. because they know the site at issue (due to its overall excellent reputation, self-evident trustworthiness, etc.) and would have made purchases even without a certification. For such sites, the left side may be smaller than the right, making certification unprofitable.

This model yields a testable prediction: A negative relationship between site popularity and site certification, i.e. that less popular sites are more likely to seek certification. Well-known sites have small $\mathbf{r_i}$'s because users already know them. For example, a certification probably cannot boost eBay's reputation with typical consumers; a certification gives little additional information about eBay's trustworthiness. In contrast, smaller sites have bigger $\mathbf{r_i}$'s because certifications are more likely to increase consumers' evaluation of site trustworthiness.

This model also allows extensions in the approach of Tang et al. Tang considers heterogeneity in sites' costs of protecting users' privacy or, more generally, complying with certification rules. For example, cost heterogeneity might arise because some sites' businesses entail prohibited or borderline business practices. Any such costs could be added into a given site's site-specific $\mathbf{c_i}$, explaining why some highly-untrustworthy sites remain uncertified. Conversely, some trustworthy sites might have high effective $\mathbf{c_i}$ because their complex operations or high-paid staff (e.g. outside attorneys) would make confirming compliance particularly onerous. A sufficiently large $\mathbf{c_i}$ could explain trustworthy sites nonetheless declining to obtain certification.

Note that this is a static model; it does not develop and cannot predict equilibrium outcomes. In particular, in this simple model, consumers do not update their beliefs according to sites' behavior, nor do sites change their behavior to suit changes in consumers' decision-making processes (since consumers' decisions do not change). I offer this model to help understand observed outcomes – that some site get certified and others do not. I omit a more general model because I consider this market far from equilibrium (many new users arriving, many new sites appearing, widespread hidden information about site types) and because such a model is not necessary to understand or interpret the subsequent empirical results.

2.3. Adverse Selection in Search Engine Results

The empirical economics literature confirms a worry of adverse selection in search engine advertising. Animesh et al. (2005) seek to test the theory that higher quality firms tend to advertise more (e.g. by placing higher PPC bids). Following Darby and Karny (1973), Animesh et al. separate search terms according to product type: They distinguish between search goods (with characteristics identifiable through pre-purchase inspection), experience goods (with characteristics revealed only through consumption), and credence goods (for which even ex-post observation does not reveal all characteristics). For experience and credence goods, Animesh et al. find that lower quality firms bid higher, but they find a positive relationship between quality and bids for search goods. Animesh et al. therefore find an adverse selection effect in search engine advertising for experience and credence goods, though not for search goods.

Animesh et al. consider the *intensive* margin of search engine advertising – *how much* a site bids for PPC ads. Animesh et al. therefore effectively test the hypothesis of higher-ranked PPC sites being safer than lower-ranked sites. But adverse selection can also present itself at the *extensive* margin – whether sites advertise through PPC *at all*. This extensive analysis is the focus of my subsequent analysis.

In contrast, search engines' organic results use a mechanism intended to select for high-quality sites. As described in Google's much-cited PageRank specification (Brin and Page 1998), modern search engines generally evaluate sites in part based on their inbound links, i.e. links from other sites. A "bad" site is likely to find it harder to obtain inbound links: Other sites won't want to link to a site that they consider untrustworthy. So link-based rating systems may make search engines' organic listings more trustworthy and less subject to adverse selection or manipulation.

A testable implication arises from the theory of organic search as a defense to adverse selection: Higher-quality search engines should offer safer organic search results than those at lower-quality search engines. Here, quality refers most naturally to use of PageRank-type ratings, but more loosely to user appraisal of search engine relevance. Industry sources indicate that Google and Yahoo increasingly have comparable organic search quality, with Microsoft and Ask Jeeves somewhat behind, all in that order. See e.g. Webmasterbrain (2006).

Of course not all analysts believe search engine advertising is subject to an adverse selection problem. Critiquing the problem of low-quality results in search engines' organic search results, Overture founder Bill Gross reportedly commented that "the best way to clean up search results was to use money as a filter." (Hansell 2001) Gross's comment effectively asks what distinguishes high-quality sites from low-quality sites. In Gross's view, the difference that only high-quality sites can afford to advertise. But the preceding paragraphs offer an alternative theory: That low-quality sites may be equally (or better) able to advertise, but that high-quality sites can more easily obtain favorable organic search engine placement via links from other high-quality sites. My empirical strategy attempts to distinguish between these competing hypotheses.

3 Empirical Strategy

The hypotheses in Equations 8 and 9 call for comparison of the "true" trustworthiness of a large number of sites. In general this data is difficult to obtain. If such data were readily available to consumers, there would exist no hidden type problem among web sites, and there would therefore be no opportunity for adverse selection. Furthermore, in general theory suggests economists have less information than consumers – making economists ill-equipped to know facts consumers do not. Nonetheless, the peculiarities of online trust make it possible to examine, measure, and analyze sites' trustworthiness, even though consumers and trust authorities generally lack this information.

To determine sites' "true" trustworthiness, I use data from SiteAdvisor. (Disclosure: SiteAdvisor is a for-profit firm, and I serve on its Advisory Board.) In an

effort to protect consumers from unsafe web sites, SiteAdvisor has designed automated systems ("robots") to visit web sites and attempt to measure their safety. SiteAdvisor's robots measure site characteristics likely to be of particular interest to users, but generally difficult for users to discern (except through SiteAdvisor). For example, one SiteAdvisor robot provides a different single-use email address to any web form it observes, and SiteAdvisor's systems subsequently measure how many messages were sent to that address – a basic measure of the amount of email (likely spam) a user would receive after signing up at the corresponding site. Another SiteAdvisor robot downloads all programs it finds within each site it tests; the robot installs each program on a separate virtual computer, then scans each virtual computer for spyware, assessing the possibility of spyware infection at each site. Other robots check for excessive pop-up ads, security exploits, scams, links to other bad web sites, and other behaviors likely to be unappealing

SiteAdvisor's measurements are imperfectly correlated with the objectives of trust authorities. For example, a site could send hundreds of emails per week to its registrants and members, yet still receive a TRUSTe certification and still qualify to advertise at major search engines. Nonetheless, I believe SiteAdvisor's measurements are highly correlated with the behaviors users actually find objectionable. As discussed above, users are unlikely to understand the subtleties of trustworthiness certifications; rightly or wrongly, users seem to regard such certifications as general seals of approval and of good business practices. Any site tripping SiteAdvisor's robots is a site likely to be of substantial concern to typical users. I therefore consider SiteAdvisor's data a good proxy for sites' true trustworthiness – for the outcomes users actually care about.

Separately, I need data on trust authorities' member lists. I obtain member lists from the current web sites of TRUSTe and BBBOnLine, and I obtain yearly historic TRUSTe member lists from date-stamped data at archive.org.

For assessment of search engines' implicit endorsements of trustworthiness, I use a crawler to extract search engine results and ads as of January 2006. I extract data for 1,397 popular keywords, including all Google Zeitgeist 2005 keywords (popular and notable search terms) as well as corresponding lists from other search engines. I extract data from the top five search engines: Google, Yahoo, AOL, MSN, and Ask Jeeves. For each search term, I extract the top 50 organic results and up to the first 50 ads (if that many ads are available).

Despite the apparent simplicity of Equations 8 and 9, they hide considerable complexity. These equations might be taken to call for conditioning on other site characteristics – for example, comparing certified sites with other commercial sites rather than with a cross-section of commercial and non-commercial sites. My empirical strategy includes specifications with various controls, including a crude measure of site commerciality (.COM versus .ORG versus other extensions) as well as popularity (as measured by an ISP willing to share site popularity data).

Throughout, I analyze approximately half a million sites – generally the top sites according to the ISP that provided me with popularity data. In many specifications, I add

information about site popularity, again as measured via this ISP. My "traffic" data comes in rank form, so larger values counterintuitively imply *smaller* amounts of traffic.

4 Results and Discussion

4.1. Certification Authorities

I begin with the core comparison directly suggested by Equation 9. Comparing the trustworthiness of certified and uncertified sites, I obtain the results in Tables 1 and 2 for TRUSTe and BBBOnLine (privacy seal program), respectively. Notice that TRUSTe-certified sites are *less* likely to actually be trustworthy: Only 94.6% of TRUSTe-certified sites are actually trustworthy (according to SiteAdvisor's analysis), whereas 97.5% of all tested sites (and 97.5% of non-TRUSTe sites) are trustworthy. This analysis gives a basic initial confirmation of the adverse selection result posited in Section 2.

The TRUSTe adverse selection result in Table 1 holds in a regression framework that controls for additional variables. Table 3 Column 1 gives a probit estimation of the relationship between TRUSTe certification and true site trustworthiness. Column 2 adds an assessment of site traffic – addressing the worry that more popular sites are exogenously both safer and more likely to be certified. Column 3 adds basic notions of site type – dummies for .COM sites and for .ORG's. In each specification, the TRUSTe certification coefficient remains statistically significantly negative – indicating that, on the margin, TRUSTe certification is still associated with a reduction in the probability that a given site is actually trustworthy.

In Table 5, I test the suggestion that TRUSTe's negative association with trustworthiness is spurious. Some might worry: TRUSTe's members tend to operate complex web sites, and complex sites can fail SiteAdvisor's automated testing in more ways than simple (static, non-interactive) sites. So perhaps the untrustworthiness of TRUSTe's members reflects only that complex sites sign up with TRUSTe while complex sites also fail trustworthiness tests. I reject this hypothesis by restricting analysis to domains that offer downloads and/or email signup forms (two site characteristics reported in additional SiteAdvisor data). Restricting my analysis to this subset of domains, I find that TRUSTe certification remains significantly negative. So even among a comparison group of only complex sites – sites with interactivity as embodied in web forms or software downloads – TRUSTe's certification remains associated with less trustworthy sites.

Notably, Tables 2 and 4 indicate that BBBOnLine's privacy seal does not suffer from significant adverse selection. Both tables show that, unlike TRUSTe's certified sites, BBB-certified sites are actually more likely to be trustworthy than a random cross-section of sites. Industry sources provide insight into BBB's success: BBB staff tend to investigate applications in considerably greater detail, including requiring membership in a local better business bureau chapter (with associated additional requirements), whereas TRUSTe tends to rely primarily on applicants' self-assessments. Though BBB's approach offers important benefits, BBB apparently faces substantial difficulties: A backlog of applicants and a slow application approval process (in part reflecting the

additional required examinations and evaluations). Indeed, BBB's web site confirms that only 631 certificates have been issued to date. It is unclear whether BBB could scale its process to investigate orders of magnitude more sites (if such sites sought certification). If not, BBB's privacy seal probably cannot meaningfully evaluate a large and growing Internet. Section 5 expands on these differences and their policy ramifications.

Section 2 offered the additional hypothesis that certification authorities' membership should become increasingly untrustworthy over time. Table 6 and Figure 1 confirm that hypothesis. (Note that I use current trustworthiness as a proxy for historic trustworthiness, for lack of direct observation of sites' historic trustworthiness. This effectively assumes that trustworthy sites stay trustworthy, and vice versa. For the notion of trustworthiness used in this paper, I consider constant trustworthiness a reasonable approximation for the overwhelming majority of sites.)

4.2. Search Engines and Search Engine Advertising

In this section, I present similar analysis for search engines. Section 1.3 offers two distinct notions of search engine certification, which I consider in turn. Edelman and Rosenbaum (2006) provides further data and discussion.

Table 7 compares true trustworthiness of search engines' top results with search engines' lower-ranked organic results. Comparing the first three organic results with the entire first page (e.g. row 2 versus row 3 in Table 7) shows no substantial sign of adverse selection in trustworthiness statuses. Among the first page of organic results, lower-ranked entries seem to be slightly safer than top entries, but the difference is not statistically significant. See Table 8, Columns 1 and 2. Extending analysis to all results studied (the first five pages, or fifty results) finds top sites weakly significantly safer than lower sites. See Table 8, Columns 3 and 4. But the coefficient on site ranking is only weakly significant, and it is not economically significant for reasonable values of site ranking. This lack of adverse selection in organic search results is consistent with predictions from theory: As discussed in Section 2.2, search engines' standard analyses (e.g. link assessment and PageRank) discourage much of the strategic behavior that would lead to tainted results.

While peer assessment protects search engines' organic links, no similar system assures high quality among search engines' ads. Any site can buy advertising placements merely by signing up and paying the associated fees. Theory offers no clear predictions as to what type of sites will be most willing to pay for placements. But Table 7 answers the question empirically: Untrustworthy sites are overrepresented among ads at all five tested search engines. Notice that rows 5 through 8 (percent of sponsored results that are untrustworthy) are all larger than rows 1 through 4 (untrustworthiness of organic results) and also larger than the overall untrustworthiness rate of the entire sample of sites. An ANOVA test confirms that these differences are highly significant ($P<0.001$). Following the Bayesian updating analysis in Equation 9, consumers could reasonably conclude that search engine ads are needlessly risky – that they are better off sticking with the safer results in search engines' organic sections.

Table 7 shows a striking result for organic listings at Yahoo – *zero* untrustworthy sites obtaining the highest organic position at Yahoo (for any of the 1,397 keywords tested), compared to 2%+ at every other search engine. Conscious of Yahoo's history of manual organization of the web (e.g. via the hand-prepared Yahoo Directory), some industry sources suggest that Yahoo manually selects top organic sites for top keywords. Such a procedure could allow Yahoo to prepare exceptionally safe top organic results (for some known set of top keywords), albeit at a cost in Yahoo staff time.

5 Policy Implications

The markets at issue – for explicit "trust" certifications and for search engine advertising – are new and developing, subject to ongoing adjustment and redesign. In established industries, norms and fundamental product characteristics substantially limit how markets can develop. But in trust certifications and search engine advertising, straightforward institutional reforms could improve outcomes substantially.

In general, natural policy responses would encourage trust authorities to consider the externalities of their actions. If a trust authority says a site is trustworthy when it actually is not, at present the trust authority can largely ignore the consequences of its error. (Indeed, trust authorities may benefit from their errors: Certification authorities receive fees for each certification issued, and search engines get paid for showing advertisements, whereas no fees result from refusing certifications or ads.) But suitable sanctions could change trust authorities' calculation. An appropriate sanction would track and balance the harm consumers suffer from reliance on a trust authority's erroneous endorsement.

Regulators have clear mechanisms for influencing certification authorities' behavior. For example, law or regulation could require certification authorities to expend greater resources in finding improper or misleading uses of their certifications. Granting certifications without proper investigation could create liability for a certification entity (roughly on a negligence theory). Alternatively, policy could attempt to shape enforcement behavior. For example, requiring publication of complaints could encourage submissions and assure appropriate responses. Meritorious complaints could also be rewarded directly, i.e. via bounties. A particularly tough policy could impose minimum standards for any authority seeking to offer a general trustworthiness certification. While sites and certification authorities might criticize these approaches as unduly invasive, they are not unprecedented. For example, the FTC has already stated that a site can commit an unfair and deceptive trade practice when it violates its posted privacy policy. (See e.g. Gateway Learning 2004.)

Search engines present similar opportunities for policy intervention. A 2002 FTC rule requires search engines to label their sponsored links. But further FTC rules could extend this duty, as could litigation under existing statutes. For example, search engines could be required to exercise reasonable diligence in selecting and approving advertising buyers. Duties might track a basic common law notion of negligence – with increased care required when selling a large amount of advertising, when doing business with an otherwise-unknown company, and when selling ads in categories known to include many

untrustworthy advertisers. Search engines might satisfy these duties, at least in part, by offering improved procedures for users to complain about particular ads.

Although I focus on a few particular examples of self-regulation, my results fall within a broader context of regulators responding to incentives. Notice that certificate issuers get paid a fee when they issue certifications, but they get nothing for rejecting an application. Similarly, patent examiners at the US Patent Office get credit for each patent application approved – but far less credit for each application rejected. (Ravicher 2005) The notorious 1850 Fugitive Slave Act also set skewed statutory compensation – granting a judge twice as much money for sending a defendant to slavery as for finding a defendant free. Facing these incentives, it's little wonder that the PTO issues so many patents, or that Fugitive Slave Act courts enslaved so many victims. To earn users' trust, online trust authorities must do better – issuing certificates that actually reflect the true merits of the sites at issue.

To those who prefer private self-regulation over direct government intervention, the strong results at BBBOnLine inspire considerable optimism. BBB's tradition of self-regulation seems to help substantially – creating institutional protection against lax certifications, and blunting short-run incentives to issue certifications willy-nilly. BBB also benefits from its regional network of evaluators, whose geographic proximity to applicants lets them better assess applicants' trustworthiness. That said, BBB's small member list and apparent certification delays make it an unlikely solution to the problem of online privacy. The FTC's 1999 report seemed to hope for more, e.g. assessment of several orders of magnitude more web sites. BBB's separate Reliability seal offers a larger member list of some 27,000+ member sites – but with correspondingly less scrutiny performed on each member, and with members that turn out to be less trustworthy according to my regression analysis. The relative untrustworthiness of BBB's Reliability members suggests that BBB's Privacy seal might not scale to evaluate and certify a larger number of sites.

What lessons to take away? I advise regulators: Don't count on self-regulatory bodies to assess would-be members correctly; self-regulators' incentives diverge substantially from a reasonable social utility function. To users, I warn: Don't believe everything you read, and don't believe every trust certificate you see; those certificates may not be worth the paper they're (not) printed on. And to certification authorities: Your credibility is on the line. The world is watching.

6 Tables and Figures

*Throughout, *** denotes P-values less than 0.001, while ** denotes P-values less than 0.01 and * denotes P-values less than 0.05.*

6.1. Conditional Probability Analysis of Trust Certifications

These tables reflect analysis of top sites, as reported by a major US ISP based on its customers' web usage.

	TRUSTe-certified	Not certified
Trustworthy	874	515,268
Not Trustworthy	50	13,148

Table 1: Trustworthiness by TRUSTe Certification Status

Associated conditional probabilities:

P(trustworthy|certified) = 94.6% P(untrustworthy|certified) = 5.4%
P(trustworthy|uncertified) = 97.5% P(untrustworthy|uncertified) = 2.5%

	BBB-certified	Not certified
Trustworthy	284	515,898
Not Trustworthy	3	13,196

Table 2: Trustworthiness by BBB Privacy Certification Status

Associated conditional probabilities:

P(trustworthy|certified) = 99.0% P(untrustworthy|certified) = 1.0%
P(trustworthy|uncertified) = 97.0% P(untrustworthy|uncertified) = 3.0%

6.2. Regression Analysis of Trust Certifications

These tables reflect analysis of top sites, as reported by a major US ISP based on its customers' web usage.

Φ(Site Trustworthiness)	(1)	(2)	(3)
Constant	1.96***	1.89***	1.96***
	(0.003)	(0.005)	(0.011)
TRUSTe Certification	-0.356***	-0.302***	-0.276***
	(0.068)	(0.080)	(0.068)
Site Traffic Rank		$1.30x10^{-7}$***	$1.30x10^{-7}$***
		($6.24x10^{-9}$)	($6.24x10^{-9}$)
Site Type Dummies			Yes

Table 3: Probit of Site Trustworthiness on TRUSTe Certification and Site Characteristics

Φ(Site Trustworthiness)	(1)	(2)	(3)
Constant	1.96***	1.89***	1.96***
	(0.004)	(0.005)	(0.011)
BBB Privacy Certification	0.349	0.395	0.416
	(0.217)	(0.217)	(0.217)
Site Traffic Rank		$1.32x10^{-7}$***	$1.31x10^{-7}$***
		($6.25x10^{-9}$)	($6.25x10^{-9}$)
Site Type Dummies			Yes

Table 4: Probit of Site Trustworthiness on BBB Site Certification and Site Characteristics

Φ(Site Trustworthiness)	(1)	(2)
Constant	1.67***	1.67***
	(0.002)	(0.002)
TRUSTe Certification	-0.187*	
	(0.074)	
BBB Privacy Certification		-0.439
		(0.236)
Site Traffic Rank	$9.40x10^{-8}$***	$9.52x10^{-8}$***
	($1.00x10^{-8}$)	($1.00x10^{-8}$)
Site Type Dummies	Yes	Yes

Table 5: Probit of Site Trustworthiness on Site Certification and Site Characteristics, Among Complex Sites (with web forms and/or software downloads)

6.3. Historical Analysis of Trust Certifications

Date	Num. TRUSTe-Certified Sites	% Untrustworthy
January 1998	28	0.00%
July 1998	61	0.00%
January 1999	319	2.19%
May 1999	430	1.63%
January 2000	1467	1.77%
August 2000	1527	1.70%
January 2001	1550	2.26%
January 2002	1532	2.61%
January 2003	1208	2.24%
April 2004	1225	2.29%
July 2004	1331	2.70%
November 2004	1172	2.99%
February 2005	1263	2.93%
April 2005	1269	3.07%
January 2006	1554	3.41%

Table 6: Historical Analysis of Trustworthiness of TRUSTe-Certified Sites

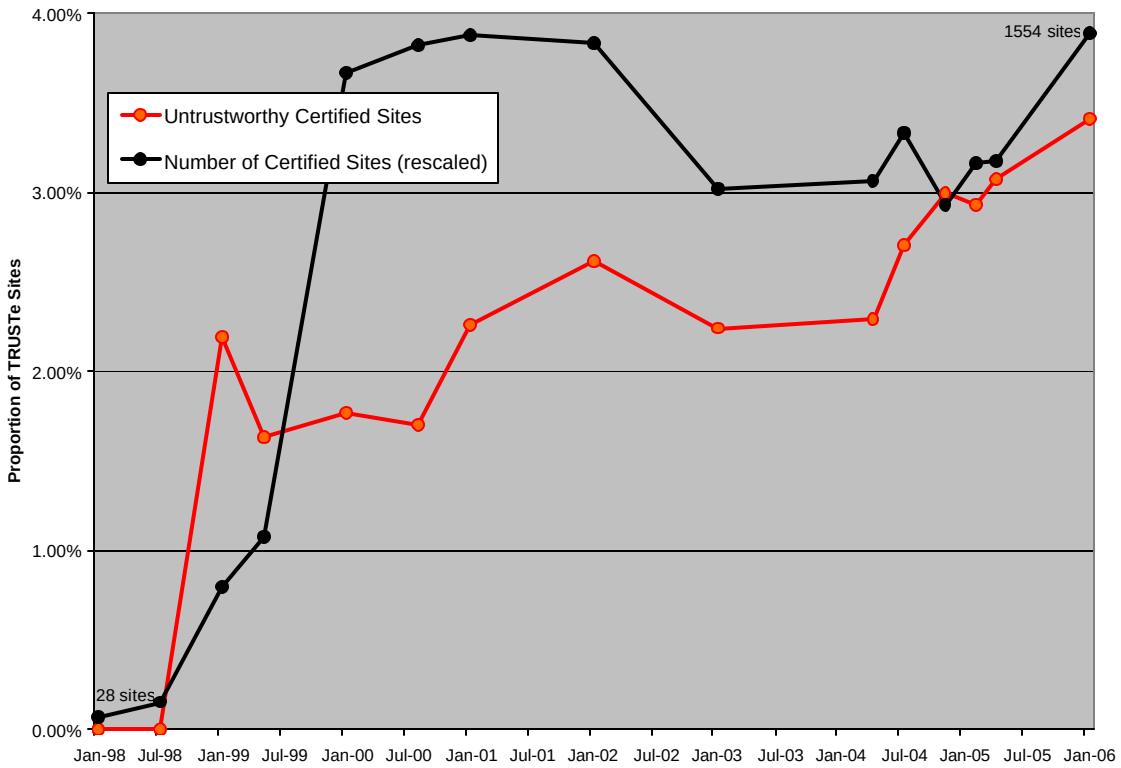

Figure 1: Historical Analysis of Trustworthiness of TRUSTe-Certified Sites

6.4. Untrustworthy Sites Certified by TRUSTe

The table below reports selected untrustworthy sites certified by TRUSTe, and a general sense of the sites' respective practices. All listed sites were certified by TRUSTe as of January 2006.

Domain	Description
Direct-revenue.com	Direct Revenue makes advertising software known to become installed without consent. Tracks what web sites users visit, and shows pop-up ads. Historically, blocks many attempts at removal, automatically reinstalls itself, and deletes certain other programs from users' PCs. Faces litigation by the New York Attorney General, plus multiple consumer class actions.
Funwebproducts.com	This site, among other Ask Jeeves toolbar distributors, installs a toolbar into users' web browsers when users agree to install smileys, screensavers, cursors, or other trinkets. Moves a user's Address Bar to the right side of the browser, such that typing an address into the standard top-left box performs a search rather than a direct navigation.
Maxmoolah.com	Offers users "free" gifts if they complete numerous sequential partner offers. Privacy policy allows sharing of user' email addresses and other information with third parties. In testing, providing an email address to Maxmoolah.com yielded a total of 485 distinct e-mails per week, from a wide variety of senders.
Webhancer.com	Makes online tracking software, sometimes installed without consent. Monitors what web sites users visit, and sends this information to Webhancer's servers.

6.5. Search Engines and Search Engine Advertising

These tables reflect analysis of a sample of 1,397 popular keywords obtained from industry sources.

	% Untrustworthy				
Which Result	Google	Yahoo	MSN	AOL	Ask
Top 1 Organic	2.73%	0.00%	2.03%	2.75%	3.23%
Top 3 Organic	2.93%	0.35%	2.24%	2.73%	3.24%
Top 10 Organic	2.74%	1.47%	2.56%	2.56%	2.94%
Top 50 Organic	3.04%	1.55%	2.46%	2.79%	3.12%
Top 1 Sponsored	4.44%	6.35%	6.17%	6.87%	7.99%
Top 3 Sponsored	5.33%	5.72%	6.16%	6.87%	7.99%
Top 10 Sponsored	5.89%	5.14%	6.37%	6.35%	8.31%
Top 50 Sponsored	5.93%	5.40%	6.01%	7.20%	8.20%

Table 7: Site Trustworthiness by Search Engine Placement Time and Position

Φ(Site Trustworthiness)	(1)	(2)	(3)	(4)
Constant	1.800***	1.844***	1.935***	1.888***
	(0.0164)	(0.025)	(0.010)	(0.013)
Organic Ranking Position	0.0153***	0.0039	-0.0059***	-0.0047***
	(0.0022)	(0.0025)	(0.0005)	(0.0005)
Search Engine Dummies		Yes		Yes
Result Restriction	Top 10	Top 10	All	All

Table 8: Probit of Site Trustworthiness on Organic Search Engine Ranking

7 Appendix: Reversibility of Conditionals in Bayes Rule Analysis, when Outcome and Signal are Both Binary

The body of the paper claims that, in the case in which both s and t are binary, $P(s|t)<P(s|\bar{}t)$ if and only if $P(t|s)<P(t|\bar{}s)$. This section provides the proof.

For s and t binary, there are four possible combinations of values of s and t. Let the values within the table below denote the respective probabilities, with a+b+c+d=1 by completeness.

	s	-s
t	a	b
-t	c	d

The definition of conditional probability yields the following identities:

$$P(s|t) = \frac{a}{a+b} \tag{14}$$

$$P(s|-t) = \frac{c}{c+d} \tag{15}$$

$$P(t|s) = \frac{a}{a+c} \tag{16}$$

$$P(t|-s) = \frac{b}{b+d} \tag{17}$$

Suppose $P(s|t)<P(s|-t)$. Substituting, then cross-multiplying and expanding:

$$\frac{a}{a+b} < \frac{c}{c+d} \tag{18}$$

$$a(c+d) < c(a+b) \tag{19}$$

$$ac + ad < ac + bc \tag{20}$$

Subtracting ac from each side:

$$ad < bc \tag{21}$$

Adding ab to each side yields:

$$ab + ad < ab + bc \tag{22}$$

$$a(b+d) < b(a+c) \tag{23}$$

$$\frac{a}{a+c} < \frac{b}{b+d} \tag{24}$$

Substituting, using the identities in (16) and (17):

$$P(t|s) < P(t|\bar{}s) \tag{25}$$

So $P(s|t)<P(s|\bar{}t)$ ➔ $P(t|s) < P(t|\bar{}s)$. But all steps in the analysis are reversible, which proves the converse and completes the proof.

8 References

Akerlof, George, 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*. **84** 488-500.

Animesh, Animesh, Ramachandran, Vandana and Viswanathan, Siva, 2005. "Quality Uncertainty and Adverse Selection in Sponsored Search Markets." .NET Institute Working Paper No. 05-27 http://ssrn.com/abstract=851286.

Baye, Michael, and Morgan, John, 2003. "Red Queen Pricing Effects in E-Retail Markets." Mimeo, http://faculty.haas.berkeley.edu/rjmorgan/Red%20Queen.pdf.

Boutin, Paul, 2002. "Just How Trusty is Truste?" *Wired*. April 9, 2002. http://www.wired.com/news/exec/0,51624-0.html.

Brin, Sergey, and Page, Lawrence, 1998. "The Anatomy of a Large-Scale Hypertextual Web Search Engine." *Computer Networks and ISDN Systems.* **30** 107-117.

Chiappori, Pierre-Andre, and Salanie, Bernard, 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy.* **108** 56-78.

Consumer Reports Web Watch. "A Matter of Trust: What Users Want From Web Sites." 2002. http://www.consumerwebwatch.org/pdfs/a-matter-of-trust.pdf.

Darby, Michael and Karny, Edi, 1973. "Free Competition and the Optimal Amount of Fraud." *Journal of Law and Economics*. **16** 67-88.

Edelman, Benjamin, and Rosenbaum, Hannah, 2006. "The Safety of Internet Search Engines." http://www.siteadvisor.com/studies/search_safety_may2006.html.

Electronic Frontier Foundation, 1999. Letter to the FTC. Original at http://www.eff.org/pub/Privacy/Email_Internet_Web/19991020_req_to_prtc_com3.html, quoted in relevant part at http://yro.slashdot.org/article.pl?sid=99/11/05/1021214.

Finkelstein, Amy, and Poterba, James, 2004. "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market." http://www.nber.org/~afinkels/papers/AFJPJPE.pdf.

FTC, 2002. "Re: Complaint Requesting Investigation of Various Internet Search Engine Companies for Paid Placement and Paid Inclusion Programs." http://www.keytlaw.com/FTC/Rules/paidplacement.htm.

In the Manner of Gateway Learning Corp., FTC File No. 042-3047. http://www.ftc.gov/os/caselist/0423047/0423047.htm.

Gaudeul, Alexandre, 2004. "Internet Intermediaries' Editorial Content Quality." Industrial Organization, Economics Working Paper Archive at WUSTL. http://econwpa.wustl.edu:80/eps/io/papers/0409/0409005.pdf.

Genesove, David, 1993. "Adverse Selection in the Wholesale Used Car Market." *Journal of Political Economy.* **101** 644-665.

Greenstadt, Rachel and Smith, Michael, 2005. "Protecting Personal Information: Obstacles and Directions." http://infosecon.net/workshop/pdf/48.pdf.

Hansell, Saul, 2001. "Paid Placement is Catching On in Web Searches." *New York Times*. June 4, 2001.

Joachims, Thorsten, Granka, Laura, Pan, Bing, Hembrooke, Helene, and Gay, Geri, 2005. "Accurately Interpreting Clickthrough Data as Implicit Feedback." *Proceedings of the Conference on Research and Development in Information Retrieval*. http://www.cs.cornell.edu/People/tj/publications/joachims_etal_05a.pdf.

LaRose, Robert, and Rifon, Nora, 2002. "Your Privacy Is Assured – Of Being Invaded: Web Sites with and without Privacy Seals." http://www.msu.edu/~larose/es2003post.htm.

Marable, Leslie, 2003. "Consumer Reaction to Learning the Truth About How Search Engines Work." http://www.consumerwebwatch.org/pdfs/false-oracles.pdf.

Ravicher, Daniel, 2005. "Statement Before the Subcommittee on Courts, the Internet, and Intellectual Property." http://www.pubpat.org/Ravicher%20Statement%20on%20Patent%20Act%20of%202005.pdf.

"Realty Tracker – TRUSTe Case Study." TRUSTe. http://www.truste.org/pdf/Realty_Tracker_Case_Study.pdf.

Singel, Ryan, 2006. "'Free IPod' Takes Privacy Toll." *Wired*. March 16, 2006. http://www.wired.com/news/technology/0,70420-0.html.

"Self-Regulation and Privacy Online." FTC. July 1999. http://www.ftc.gov/os/1999/07/privacy99.pdf.

Tang, Zhulei, Hu, Yu, and Smith, Michael, 2005. "Protecting Online Privacy: Self-Regulation, Mandatory Standards, or Caveat Emptor." http://infosecon.net/workshop/pdf/31.pdf.

"TRUSTe Fact Sheet." http://www.truste.org/about/fact_sheet.php.

"TRUSTe Program Requirements." http://www.truste.org/requirements.php.

"TRUSTe Watchdog Reports." https://www.truste.org/consumers/watchdog_reports.php.

"US Patent and Trademark Office FY 2005 Fee Schedule." http://www.uspto.gov/web/offices/ac/qs/ope/fee2004dec08.htm.

Webmasterbrain Search Engine Experiment, 2006. http://www.webmasterbrain.com/seo-tools/seo-experiments/the-search-engine-experiment/test-results/.