Electronic Commerce Research and Applications xxx (2010) xxx–xxx

Contents lists available at ScienceDirect

Electronic Commerce Research and Applications

journal homepage: www.elsevier.com/locate/ecra

Adverse selection in online "trust" certifications and search results

Benjamin Edelman

Harvard Business School, 1 Soldiers Field Rd., Boston, MA 02163, United States

ARTICLE INFO

Article history:
Received 19 October 2009
Received in revised form 19 May 2010
Accepted 23 June 2010
Available online xxxx

Keywords:
Adverse selection
Certification
Reputation
Trust
Regulation

ABSTRACT

Widely-used online "trust" authorities issue certifications without substantial verification of recipients' actual trustworthiness. This lax approach gives rise to adverse selection: the sites that seek and obtain trust certifications are actually less trustworthy than others. Using an original dataset on web site safety, I demonstrate that sites certified by the best-known authority, TRUSTe, are more than twice as likely to be untrustworthy as uncertified sites. This difference remains statistically and economically significant when restricted to "complex" commercial sites. Meanwhile, search engines create an implied endorsement in their selection of ads for display, but I show that search engine advertisements tend to be less safe than the corresponding organic listings.

1. Introduction

When agents have hidden types, contract theory warns of bad results and potentially even market unraveling. Since Akerlof's "lemons" (1970), others have worried about similar problems in markets with hidden types – like bad drivers wanting more car insurance than good drivers (Chiappori and Salanie 2000), and healthy people disproportionately buying annuities (Finkelstein and Poterba 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 research manages to address these problems. For example, Chiappori and Salanie (2000) focus on novice drivers, who have less private information about their own type (since they have not yet started to drive), letting economists 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 might affect other drivers.

This paper applies an adverse selection model to a new market: web sites and their associated "trust"-type certifications. With a proprietary data source, I analyze characteristics generally unobservable both to consumers and to trust authorities. Unmasking sites' otherwise-hidden types provides an unusual opportunity to measure the magnitude of adverse selection occurring in this market.[1]

Beyond adverse selection, trust certifications are also of interest in their own right. These 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 adverse selection significantly taints trust certifications: my analysis indicates that low-quality sites disproportionately seek and receive certification, substantially reducing overall certification quality. In particular, in Section 6, I find that sites certified by the best-known authority, TRUSTe, are more than twice as likely to be untrustworthy as uncertified sites.

Distinct from trust certifications, search engines effectively endorse certain sites through their selection of ads to present with users' search results. Seeing prominent advertisements juxtaposed with sites presented as the "most relevant" for a given search, users may mistakenly believe the advertised sites deserve their trust. But in fact the true test for appearance in search advertisements is paying a fee, not satisfying substantive requirements. In Section 7, I show that search engine advertisements are systematically less safe than the corresponding organic results.

2. The basic web site safety problem

Consumers seeking online services face a serious difficulty in deciding what sites to use. Consumers could stick with "known-good" big names, but such a narrow focus would reduce match quality, denying users the rich diversity of Internet content. Exploring the broader Internet offers the potential for a better match, but with important risks: untrustworthy sites might send users spam (if users register or otherwise provide email addresses), infect users' computers with viruses or other harmful code (if users install the programs that sites offer), or simply fail to deliver the

E-mail address: bedelman@hbs.edu

[1] Despite similarity in name, trust certification authorities are entirely distinct from the certification authorities that offer SSL certificates, code-signing certificates, and encryption keys for secure communications.

1567-4223/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.elerap.2010.06.001

Please cite this article in press as: Edelman, B. Adverse selection in online "trust" certifications and search results. *Electron. Comm. Res. Appl.* (2010), doi:10.1016/j.elerap.2010.06.001

promised merchandise (if users make purchases). Ex ante, users have no easy way to know which sites to trust. A safe-looking 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 their integrity. But as usual in adverse selection, bad sites pretend they are good.

Facing numerous untrustworthy or even malicious sites, some analysts call for government regulation. In principle, a government agency might examine web sites in search of spam, scams, and harmful programs. To some extent, the FTC and state attorneys general perform such investigations – though their efforts address only a small portion of bad actors. As a practical matter, government intervention seems inapt. For example, Tang et al. (2005) present a model of enforcement of online privacy breaches, finding mandatory government standards appropriate only for the most serious harms.

At the other extreme, users might be left entirely on their own. In complete *caveat emptor* ("buyer beware"), no regulator, computer maker, or IT department helps cure a user's problems. In some respects, *caveat emptor* is a reasonable description of the current state of affairs. (IT departments cannot protect users from getting ripped off, and even computer experts often feel powerless to stop spam.) But unaccountability carries substantial costs – leading users to take excessive precautions, and 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 of fraud and abuse.

Finally, there exists a middle approach between the extremes of government regulation and *caveat emptor*: a non-governmental rating organization. Such an organization would identify specific bad practices, then evaluate sites' behaviors. If evaluations were accurate and low-cost, such ratings might support an equilibrium where good firms receive positive evaluations, and where consumers use only sites with positive ratings. Tang et al. (2005) suggest that rating organizations are appropriate for a broad class of online interactions.

3. Trust authorities

Most prominent among non-governmental rating organizations are so-called "trust" certification authorities. These organizations set out specific criteria for membership, often focusing on privacy or on online safety more generally. The organizations reward their members by offering seals to be placed on recipients' web sites, typically on registration forms and checkout pages. To date, the best-known trust authorities are TRUSTe and BBBonline.

In principle, trust 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, critics question the effectiveness of certain trust authorities. LaRose and Rifon (2002) offer a stinging critique: trust authorities have granted multiple certifications to firms under investigation by the FTC for privacy policy violations; trust authorities have declined to pursue complaints against major companies whose privacy breaches were found to be "inadvertent"; and in one case a trust authority even failed to abide by *its own* privacy policy. Ryan (2006) raises similar concerns: in a 2004 investigation after user complaints, TRUSTe gave Gratis Internet a clean bill of health. Yet subsequent New York Attorney General litigation uncovered Gratis' exceptionally far-reaching privacy policy violations – selling 7.2 million users' names, email addresses, street addresses, and phone numbers, despite a privacy policy exactly to the contrary.

As a threshold matter, trust authorities' substantive standards often seem to duplicate existing duties or practices. Consider the obligations in TRUSTe's Program Requirements. The first listed rule, requiring an email unsubscribe function, duplicates Sec.5.(a)(4)(A) of the federal CAN-SPAM Act. Similarly, credit card network rules exactly overlap with TRUSTe's requirement of SSL encryption (or similar technology) to protect sensitive credit card numbers. Boutin (2002) reports that TRUSTe initially lacked any substantive requirements whatsoever (requiring only the *presence* of a privacy policy). Low standards match the predictions of Lizzeri (1999), finding that, under general conditions, a certification intermediary prefers only to reveal whether quality exceeds some minimal standard.

Tellingly, strikingly few certificates have been revoked. For example, the TRUSTe Fact Sheet reports only two certifications revoked in TRUSTe's 10-year history. TRUSTe's small staff has little apparent ability to detect infractions. Instead, TRUSTe's posted procedures emphasize user complaints and sites' self-certifications. When violations have been uncovered, the proof has come from outside complaints, not from TRUSTe itself.

TRUSTe's "Watchdog Reports" also indicate a lack of focus on enforcement. TRUSTe's postings reveal that users continue to submit hundreds of complaints each month. But of the 3416 complaints received since January 2003, TRUSTe concluded that *not a single one* required any change to any 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, trust authorities are paid by the same companies they certify; in the language of Greenstadt and Smith (2005), trust authorities are "captured". With this revenue model, authorities have little short-run incentive to seek higher standards: any such pressure would discourage renewals and future applications – reducing revenues.

Even the creators of trust authorities report disappointment in their development. TRUSTe co-founder Esther Dyson called TRUSTe "a little too corporate", and said TRUSTe lacks the "moral courage" to criticize violations (Boutin 2002). Similarly, the Electronic Frontier Foundation, another TRUSTe co-founder, told the FTC that "it is time to move away from a strict self-regulation approach" (1999).

Table 1 reports selected untrustworthy sites certified by TRUSTe, along with a general statement of the sites' respective practices. As of January 2006, TRUSTe listed all these sites among its certified members.

Facing allegations of low substantive standards, lax enforcement, and ethical compromise, it is unclear what direct benefits site certifications offer to consumers. But at least some consumers seem to regard certification as a significant positive signal. For example, in recruiting web sites to get certified, TRUSTe offers an endorsement from certificate recipient Realty Tracker, which says TRUSTe "convey[ed] trust" and "built confidence" with site visitors, yielding "an increase in registrations". See TRUSTE's Realty Tracker Case Study.

Moreover, firms are well-equipped to evaluate claimed benefits to certification: firms can randomly include or omit a seal, thereby measuring whether a seal increases registrations and sales. Indeed, year after year, hundreds of firms seek and renew TRUSTe certification – suggesting that firms find certification valuable. Furthermore, in the related context of comparison shopping sites, Baye and Morgan (2003) empirically confirm the benefits of certification: merchants with seals can charge a price premium without losing customers.

Even well-known web sites tout their safety certifications. For example, the Microsoft's Online Privacy Policy index features the

[2] For example, TRUSTe failed to update its Watchdog Reports list between June 2004 and spring 2006, an omission corrected only after circulation of a draft of this article. Even in 2009, Watchdog Reports suffer broken links, missing reports, and contradictory document titles.

Please cite this article in press as: Edelman, B. Adverse selection in online "trust" certifications and search results. *Electron. Comm. Res. Appl.* (2010), doi:10.1016/j.elerap.2010.06.001

Table 1
Selected untrustworthy sites certified by TRUSTe.

Domain	Description
Direct-revenue.com	Makes advertising software known to become installed without consent. Tracks what web sites users visit, and shows pop-up ads. Blocks many attempts at removal, and automatically reinstalls itself, including by disguising itself as a printer driver. Deletes competing advertising software from users' PCs. Faced litigation by the FTC and the New York Attorney General, plus multiple consumer class actions
Funwebproducts.com	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. Shows seven sponsored links above the first organic result – overwhelming users with ads
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 observed becoming installed without user consent. Monitors what web sites users visit, and sends this information to Webhancer's servers

TRUSTe name and logo adjacent to the page's title and Microsoft logo. eBay presents its TRUSTe certification on its main registration page (a necessary step for all new users joining eBay).

Whatever the actual merits of certification authorities as arbiters of trust, some government authorities seem to regard these organizations as an appropriate step forward. For example, the FTC's "Self-Regulation and Privacy Online" (1999) endorsed private-sector trust authorities as an alternative to comprehensive regulation of online privacy and safety.

The FTC's 1999 recommendation specifically cites two well-known certification systems: TRUSTe's Web Privacy Seal and BBBOnLine's Privacy Seal. My subsequent analysis focuses on these authorities due to their prevalence, their relatively large member lists, and the public availability of their member lists.

4. Theory of adverse selection in trust authorities

Suppose certain trust authorities issue certifications of trustworthiness without rigorous assessment of recipients' true trustworthiness. Certifications seek to signal consumers that the certified firms are in fact highly likely to be trustworthy. But if untrustworthy firms can easily get certified, the signal drops in value: seeing a certified firm, a consumer would rightly worry that the firm is not truly trustworthy.

To provide a positive signal, a certification must increase a rational consumer's assessed probability that a site is trustworthy. Suppose a rational consumer has a prior belief $P(t)$ that a given site is trustworthy. The consumer then receives a signal ("s") of trustworthiness ("t"). The consumer updates according to Bayes Rule:

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

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 $P(t|s) > P(t)$, which implies:

$$\frac{P(s|t)}{P(s|t)P(t) + P(s|\bar{t})P(\bar{t})} > 1 \tag{3}$$

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

$$P(s|t) > P(s|\bar{t}) \tag{4}$$

Eq. (4) offers an intuitive result: for a certification to increase a consumer's assessment of the probability that a certified site is safe, the certification must be given to trustworthy sites more often than it is given to untrustworthy sites.

4.1. Testing for adverse selection at trust authorities

Eq. (4) 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.) Then an informative certification requires:

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

Adverse selection offers a clear empirical prediction: that the inequality in (5) should fail. In particular, if adverse selection substantially affects these certifications, then certified sites should be *less* safe than uncertified sites.

Hypothesis 1: Certified sites are less safe than uncertified sites.

Analyzing correlations between trustworthiness and certification continues the approach in the adverse selection literature. Consider Finkelstein and Poterba (2004), finding that annuitants live longer than non-annuitants – a negative relationship between claimed type (annuity purchase) and outcome (lifetime). Chiappori and Salanie (2000) uses a similar 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). Genesove (1993) extends these correlations with the equilibrium assumption that average price in a given market must reflect average quality in that market. Genesove 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), interpreting a significant coefficient as evidence of adverse selection at used car dealers. Villeneuve (2003) offers a specific measurement of "intensity of adverse selection", calculated as the quotient between the prevalence of some action (e.g. buying insurance) in a subsample, versus the action's prevalence in the full population. Rearranging terms, Villeneuve's measure matches (5).

Others studying online trust authorities have also worried of adverse selection. For example, LaRose and Rifon (2002) find that privacy policies at certified sites allow more invasive data collection than policies at uncertified sites. But where LaRose and Rifon hand-score 200 sites, I use automation to evaluate hundreds of thousands of sites, and I consider axes of trustworthiness other than privacy policy loopholes. In addition to investigating the quality of certified sites, Jamal et al. (2003) specifically consider certifiers lowering their standards to attract more sites. But Jamal et al. study only 34 well-known sites certified as of 2001 – limiting the generality of their findings. In contrast, my analysis includes more recent data and far more sites.

Please cite this article in press as: Edelman, B. Adverse selection in online "trust" certifications and search results. *Electron. Comm. Res. Appl.* (2010), doi:10.1016/j.elerap.2010.06.001

4.2. Trust authorities in equilibrium

Critics might reasonably doubt whether uninformative certifications 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 actually uninformative. Then users should discount or ignore those certifications. But if consumers ignore the certifications, sites have no incentive to become certified. Then certification schemes should disappear altogether.

It is reassuring to predict that worthless trust authorities will collapse. But as an empirical matter, trust authorities have existed for some time and show no sign of disappearing. Although inconsistent with a world of fully-informed consumers, the persistence of trust authorities makes sense under reasonable assumptions. For example, suppose some users are slow learners – drawing inference about certification based on the quality of sites certified *in prior periods*. Then an initial batch of high-quality certified sites would effectively subsidize future certifications.[3] Alternatively, if some users are naïve (mistakenly trusting certifications that are actually worthless), certification would be profitable if naïve users are sufficiently widespread relative to the cost of certification. In extensions (available on request), I have sketched a model of these effects.

The slow-learner model of user behavior yields an empirical prediction: The average quality of certified sites should decrease over time. Suppose a trust authority happened to start with members that truly are trustworthy, producing a favorable initial reputation with users. (Consider the alternative: If a certifier began by certifying untrustworthy sites, it would have little hope of building a positive reputation with users.) In the face of slow learning, that favorable reputation would take some time to dissipate. In the interim, untrustworthy firms can profit from certification. The resulting hypothesis:

> Hypothesis 2: Trust authorities do not suffer adverse selection in initial periods, but they suffer adverse selection that worsens over time.

5. Empirical strategy

The preceding hypotheses call for analysis of true trustworthiness of a large number of sites. In general this data is difficult to obtain. If consumers knew sites' actual trustworthiness, there would be no hidden types and no opportunity for adverse selection. But new data collection systems allow analysis of sites' actual behaviors even though consumers and trust authorities largely 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.) To protect consumers from unsafe web sites, SiteAdvisor runs automated systems to visit web sites and attempt to measure their safety. SiteAdvisor's automation uncovers site characteristics that are otherwise difficult for users to discern. For example, one SiteAdvisor system provides a different single-use email address to each web form it finds. SiteAdvisor measures how many messages are subsequently sent to that address – identifying sites and forms that yield junk mail. Another SiteAdvisor system downloads all programs it finds, installs each program on a separate virtual computer, then scans for spyware – assessing the possibility of infection at each site. Other systems check for excessive pop-ups, security exploits, scams, links to other bad sites, and more. Typically, a site fails SiteAdvisor's evaluation if it fails any of these checks.

SiteAdvisor's measurements are imperfectly correlated with trust authorities' stated rules. For example, a site could send its registrants hundreds of e-mails per week, yet still receive a TRUSTe certification. Nonetheless, SiteAdvisor's stated approach is highly correlated with the behaviors *users* deem objectionable. Without understanding the subtleties of trust authorities' rules, users seem to regard certifications as general indicators of good business practices; the very words "trust" (in TRUSTe) and "better business" indicate *overall* ethical practices, and do not suggest that good practices are limited to particular facets of a company's operations. Meanwhile, any site failing SiteAdvisor's tests is a site likely to present substantial concern to typical users; for example, a site sending enough email to trigger SiteAdvisor's alarm is a site where few users would want to register, and a site distributing malware detected by SiteAdvisor's scanners is a site where few users would want to download software. I therefore consider SiteAdvisor data a good proxy for sites' true trustworthiness – for the outcomes users actually care about, even when those outcomes differ from trust authorities' official requirements.

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 the Internet Archive (archive.org).

Table 2 presents SiteAdvisor's policies and compares these policies with the requirements of TRUSTe and BBBOnLine.

Eqs. (4) and (5) 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 full cross-section of sites. My analyses include specifications with various controls, including a crude measure of site commerciality (.COM versus .ORG versus other extensions) as well as popularity (as measured by a large US ISP).[4,5] I analyze approximately half a million sites – the web's top sites according to the ISP that provided popularity data.

6. Results and discussion

I begin by testing Hypothesis 1 using the method in Eq. (5). Comparing the trustworthiness of certified and uncertified sites (within the top web sites reported by my ISP data source), I obtain the results in Tables 3 and 4 for TRUSTe and BBBOnLine (privacy seal program), respectively.

Computing conditional probabilities from Table 3 yields the pie charts in Fig. 1. 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 data), whereas 97.5% of non-TRUSTe sites are trustworthy. That is, TRUSTe-certified sites are more than twice as likely to be untrustworthy as uncertified sites. This analysis gives an initial confirmation of the adverse selection result posited in Section 4.

The TRUSTe adverse selection result in Table 3 holds in a regression framework that controls for additional variables. In Table 5, column 1 gives probit estimation of the relationship between TRUSTe certification and true site trustworthiness. Column 2 adds site traffic – addressing the worry that popular sites are exogenously

[3] This chronology seems to match the history of TRUSTe, which was founded by a set of trustworthy companies to serve their regulatory goals. In particular, those companies preferred private-sector certification as an alternative to threatened FTC regulation of online privacy practices. Only later did TRUSTe begin to certify sites with more controversial practices.

[4] Popularity data comes in rank form, so larger values imply lesser traffic.

[5] By agreement with the ISP, I cannot identify it by name.

Please cite this article in press as: Edelman, B. Adverse selection in online "trust" certifications and search results. *Electron. Comm. Res. Appl.* (2010), doi:10.1016/j.elerap.2010.06.001

Table 2
Comparison of selected TRUSTe, BBB privacy, and SiteAdvisor policies.

Characteristic	TRUSTe web site privacy seal	BBBOnLine privacy	SiteAdvisor
Software downloads	No restriction	No restriction	Rates a site unfavorably if the site offers programs that are, or that bundle, "adware" or "spyware"
Email	No restriction	No restriction	Rates a site unfavorably if the site sends a large number of messages or does not honor requests to unsubscribe
Web links	No restriction	No restriction	Rates a site unfavorably if the site links to other sites rated unfavorably
BBB membership	No requirement	Required, with a satisfactory record of handling complaints	No requirement
Privacy policy	Compulsory. Site must self-certify its practices. Must disclose information collection and use	Compulsory. Three dozen rules about privacy policy provisions and site practices	No requirement
Dispute resolution with consumers	Site must accept consumer complaints and participate in TRUSTe "Watchdog" process	Site must participate in the BBBOnline Dispute Resolution Process	n/a
Application or certification fee	Yes, up to $7999 per year	Yes, up to $7000 per year	No

Table 3
Trustworthiness by TRUSTe certification status.

	TRUSTe-certified	Not certified
Trustworthy	874	515,309
Not Trustworthy	50	13,148

Table 4
Trustworthiness by BBB privacy certification status.

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

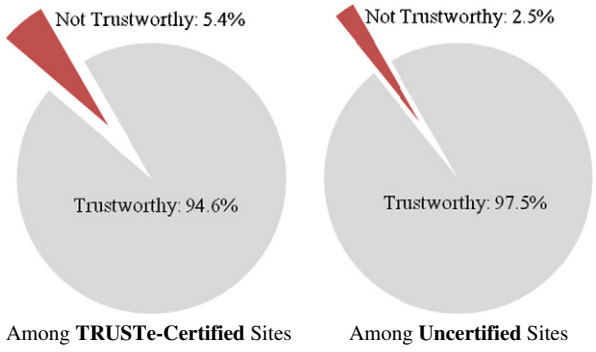

Fig. 1. Comparing TRUSTe-certified and uncertified sites.

both safer and more likely to be certified. Column 3 adds a notion of site type – dummies for .COM sites and for .ORG's. In each specification, the TRUSTe certification coefficient remains significantly negative. That is, on the margin, TRUSTe certification remains associated with a reduction in the probability that a given site is actually trustworthy.

In Table 6, Column 1, 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 both (1) get certified by TRUSTe, and (2) fail automated trustworthiness tests. I reject this hypothesis by restricting analysis to domains that offer downloads and/or email signup forms. Restricting my analysis to this subset of domains, I find that the coefficient on TRUSTe certification remains significantly negative.

Notably, BBBOnLine's privacy seal does not suffer significant adverse selection. Unlike TRUSTe's certified sites, BBB-certified sites are slightly more likely to be trustworthy than a random cross-section of sites (see Fig. 2). This result holds in a regression framework (Table 7) including when controlling for site complexity (Table 6, Column 2). Industry sources attribute BBB's success to BBB's detailed evaluation of applicants. For example, BBBOnLine requires that an applicant be a member of a local BBB chapter (which adds its own requirements), whereas TRUSTe tends to rely primarily on applicants' self-assessments. Though BBB's approach offers important benefits, BBB apparently faces substantial difficulties including a backlog of applicants and a slow application approval process (in part caused by the additional required evaluations). BBB's web site reports only 631 certificates issued to date, and it is unclear whether BBB could scale its process to evaluate orders of magnitude more sites. Section 5 expands on the policy ramifications of these differences.

Hypothesis 2 conjectured that over time, certification comes to include less trustworthy sites. Using historical TRUSTe membership data preserved by Archive.org, Table 8 and Fig. 3 confirm that hypothesis as to TRUSTe. Note that my SiteAdvisor data all dates from 2006; I do not observe sites' prior practices. Instead, I use sites' then-current trustworthiness as a proxy for historic behavior

Table 5
Probit of site trustworthiness on TRUSTe certification and site characteristics.

Φ (site trustworthiness)	(1)	(2)	(3)
Constant	1.96^{***} (0.003)	1.89^{***} (0.005)	1.96^{***} (0.011)
TRUSTe certification	-0.356^{***} (0.068)	-0.302^{***} (0.080)	-0.276^{***} (0.068)
Site traffic rank		$1.30 \times 10^{-7***}$ (6.24×10^{-9})	$1.30 \times 10^{-7***}$ (6.24×10^{-9})
Site type dummies			Yes

Throughout all regressions, $***$ denotes P-values less than 0.001, $**$ denotes P-values less than 0.01, and $*$ denotes P-values less than 0.05.

Please cite this article in press as: Edelman, B. Adverse selection in online "trust" certifications and search results. *Electron. Comm. Res. Appl.* (2010), doi:10.1016/j.elerap.2010.06.001

Table 6
Probit of site trustworthiness on site certification and site characteristics, among complex sites (with web forms and/or software downloads).

Φ (site trustworthiness)	(1)	(2)
Constant	1.67*** (0.002)	1.67*** (0.002)
TRUSTe certification	−0.187* (0.074)	
BBB privacy certification		−0.439 (0.236)
Site traffic rank	9.40×10^{-8}*** (1.00×10^{-8})	9.52×10^{-8}*** (1.00×10^{-8})
Site type dummies	Yes	Yes

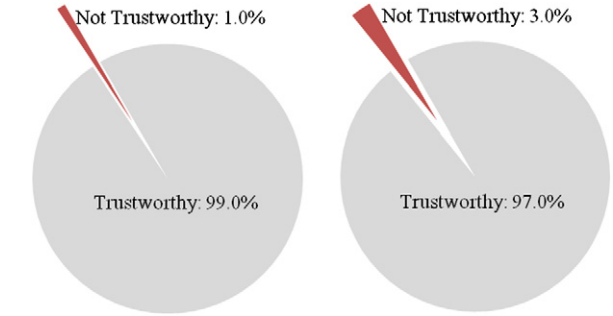

Fig. 2. Comparing BBB-certified and uncertified sites.

Table 7
Probit of site trustworthiness on BBB privacy certification and site characteristics.

Φ (Site Trustworthiness)	(1)	(2)	(3)
Constant	1.96*** (0.004)	1.89*** (0.005)	1.96*** (0.011)
BBB Privacy Certification	0.349 (0.217)	0.395 (0.217)	0.416 (0.217)
Site Traffic Rank		$1.32x10^{-7}$*** (6.25×10^{-9})	$1.31x10^{-7}$*** (6.25×10^{-9})
Site Type Dummies			Yes

Table 8
Historical trustworthiness of TRUSTe-certified sites.

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

– effectively assuming that trustworthy sites stay trustworthy, and vice versa. Based on the nature of the malfeasance giving rise to negative SiteAdvisor assessments of site behavior, and based on my discussions with SiteAdvisor management, I believe the untrustworthy sites at issue are highly likely to continue that

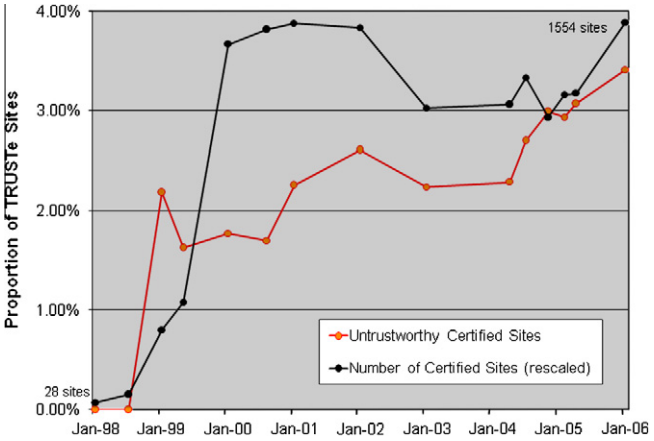

Fig. 3. Historical analysis of trustworthiness of TRUSTe-certified sites.

behavior over time, meaning that little error results from comparing historical TRUSTe membership with 2006 SiteAdvisor evaluation.

While I focus on online trust authorities certifying web site practices, other certifications seek to verify different aspects of behavior. For example, Backhouse et al. (2005a,b) examine the Certification Authorities (CAs) that issue electronic signatures for use in public key infrastructure – evaluating whether CAs issue certifications not properly justified under governing policy, and assessing the incentives that influence CAs' operations. My finding of disproportionate unwarranted certifications tracks Backhouse's prognosis of substantial quality uncertainty in public key certificates.

7. Safety of search engine advertisements

Discussion thus far has focused on explicit certifications by "trust" authorities. But other institutions also offer programs that effectively endorse certain sites. Of particular interest: In selecting which companies to feature in prominent advertisements, search engines offer placements that users may reasonably view as signaling at least tacit approval. Indeed, search engines affirmatively invite and encourage this view. For example, at Google, the "Basics" instructions to users indicate that Google will feature "the most relevant results". As to ads, Google's 8000+ word Advertising Content Policy suggests that Google diligently evaluates the advertisements it shows. Google's public remarks are in accord; when asked how Google handles deceptive ads, a spokesperson said standard policy is "we take them down" (Claburn 2006). Despite these many claims of featuring high-quality results, search engine listings can also be subject to adverse selection, wherein an identifiable subset of prominent listings is importantly less safe than other portions of the site.

7.1. Theory of adverse selection in search engine listings

High placement at a search engine implies that a search engine believes the site ranks among the best resources for a given search term – an endorsement if not a certification (Gaudeul 2004). Empirical analysis confirms that users value 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 et al. 2005). Because users tend not to understand the difference between paid search engine advertising and ordinary "organic" listings (Consumer Reports Web Watch 2002), Marable's result likely applies to all search results, not just organic results.

Please cite this article in press as: Edelman, B. Adverse selection in online "trust" certifications and search results. *Electron. Comm. Res. Appl.* (2010), doi:10.1016/j.elerap.2010.06.001

The economics literature confirms the worry of adverse selection in search engine advertising. Animesh and Ramachandran (2010) examine relationships between product type, quality (e.g. trustworthiness), and advertising strategy. For search goods (where relevant characteristics are identifiable prior to purchase), Animesh finds a positive relationship between quality and advertising bids – indicating that consumers are successfully finding high-quality providers. But for experience and credence goods (where characteristics are unknown before purchase), low-quality firms can afford to bid higher – suggesting that users cannot tell which firms are trustworthy.

Animesh et al. consider the *intensive* margin of search engine advertising – how much a site bids for search ads. Animesh therefore effectively tests the hypothesis of higher-ranked pay-per-click sites being safer than lower-ranked sites. But adverse selection can also present itself at the *extensive* margin – whether sites advertise through search advertising *at all*. In subsequent tests, I examine both of these possibilities.

In contrast to search engine ads, where bids largely determine placement, organic results are 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 (links from other sites). "Bad" sites find it harder to obtain inbound links: Others don't want to link to sites they consider untrustworthy. So link-based rating systems may make search engines' organic listings more trustworthy and less subject to adverse selection or manipulation. In particular:

Hypothesis 3: Organic results are safer than sponsored results.

Not all analysts believe search engine advertising faces adverse selection. Bill Gross, founder of pay-per-click advertising powerhouse Overture (now part of Yahoo), reportedly commented that "the best way to clean up search results [is] to use money as a filter" (Hansell 2001). Gross suggests that high-quality sites differ from low-quality sites in that only the former can afford to advertise. Hypothesis 3 suggests an alternative: that low-quality sites are equally (or better) able to advertise, but that high-quality sites can more easily obtain favorable organic placement via links from other high-quality sites. I distinguish between these theories in my test of Hypothesis 3.

Table 9
Site trustworthiness by search engine placement time and position.

Which result	% Untrustworthy				
	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

7.2. Empirical strategy and results

To test these hypotheses, I extract search engine results and ads as of January 2006. I consider 1397 popular keywords, including all Google Zeitgeist 2005 keywords (popular searches) plus similar lists from other search engines. I extract the first 50 results and up to 50 ads (if available) from each of the top five search engines, namely Google, Yahoo, AOL, Microsoft (then "MSN"), and Ask. I then check safety of each listed site using the SiteAdvisor data described in Section 5.

Table 9 demonstrates that untrustworthy sites are overrepresented among ads at all five tested search engines, relative to their presence in organic listings for the same terms: rows 5–8 (percent of sponsored results that are untrustworthy) are all larger than rows 1–4 (untrustworthiness of organic results). An ANOVA test confirms that these differences are highly significant ($P < 0.001$). These results affirm Hypothesis 3.

The relative untrustworthiness of search engine ads raises the question of why search engines' organic listings are safer. Table 10 reports that that top organic results tend to be somewhat safer than lower-ranked results – suggesting that organic listing algorithms (e.g. Google PageRank and similar) prevent untrustworthy sites from achieving high organic positions. Meanwhile, search engines viewed (per industry sources, e.g. Webmasterbrain 2006) as offering inferior organic search results also tend to feature more untrustworthy listings in top organic results. In Table 9, note the stark difference between safety of organic results at Ask with those for Google.

8. Policy implications

The framework of Akerlof (1970) offers suggestions to address problems of information asymmetry, but these responses appear to be inapt or unsuccessful in this context. To Akerlof's suggestion of *guarantees* comes the problem that, at least as currently structured, online trust authorities are in no position to offer a meaningful guarantee of certified sites' practices. Indeed, TRUSTe's Terms and Conditions specifically prohibit a user from relying on TRUSTe's certifications, and BBBOnLine's Terms of Use disclaim liability for listings. Furthermore, while a guarantee would certainly benefit users, a heightened level of verification would present certification authorities with higher costs in certification, substantial liability in case of breach by a certified site, or both. So certification authorities are unlikely to offer guarantees voluntarily. In fact, Google has intensely defended its immunity from suit if a user sees an ad among Google search results, clicks the ad, and is cheated by the advertiser. See Goodard v. Google, granting Google's Motion to Dismiss, which argued that only the advertiser, but not Google, is responsible for an advertisement's content.

Akerlof next suggests the use of *brand names* to remedy information asymmetries. To some extent "TRUSTe" and "BBB" present useful brand names that consumers can recognize and, in due course, credit or discount as appropriate. But at least in the context of TRUSTe, the value of the brand – through historical placement on well-known sites like Microsoft and eBay – becomes diluted by less trustworthy sites that later received and promoted TRUSTe

Table 10
Probit of site trustworthiness on organic search engine ranking.

Φ (site trustworthiness)	(1)	(2)	(3)	(4)
Constant	1.800*** (0.0164)	1.844*** (0.025)	1.935*** (0.010)	1.888*** (0.013)
Organic ranking position	0.0153*** (0.0022)	0.0039 (0.0025)	−0.0059*** (0.0005)	−0.0047*** (0.0005)
Search engine dummies		Yes		Yes
Result restriction	Top 10	Top 10	All	All

Please cite this article in press as: Edelman, B. Adverse selection in online "trust" certifications and search results. *Electron. Comm. Res. Appl.* (2010), doi:10.1016/j.elerap.2010.06.001

certification. Akerlof presupposes that a brand name will elect to protect and preserve its reputation, but TRUSTe's certifications indicate otherwise.

Finally, Akerlof notes the possibility of *licensing*. Certainly government oversight of online trust authorities could rein in certifications too easily granted. Conceivably some middle ground could preserve a portion of the decentralization, flexibility, and cost-saving benefits of self-regulation while adding additional control through government oversight. But to those who founded online trust authorities in the spirit of self-regulation, detailed government oversight is likely to be viewed as highly undesirable.

Seeing an apparent failure by at least some well-known trust authorities, the FTC might reasonably revisit its 1999 decision to favor certification-based self-regulation in lieu of substantive FTC oversight. But if regulators sought to retain the basic approach of self-regulatory certifications, they have ample tools to improve certification outcomes.

For one, trust authorities might appropriately face liability for their most egregious misclassifications. Within the framework of Akerlof, this approach essentially comprises a compulsory regulation-mandated guarantee – but if market forces do not inspire trust authorities to provide such guarantees, regulation could assist. At present, if a trust authority says a site is trustworthy when it is not, the trust authority currently can largely ignore the consequences of its error. (Indeed, in the short-run, trust authorities benefit from such errors: certifying a site yields a fee, while no fee results from denying a certification.) But if a trust authority falls short of a reasonable standard of care, it might properly face liability on a negligence theory. (Analogous liability has been sought, with mixed results, as to erroneous certifications by rating agencies and auditing firms in the financial sector.)

In a narrower change, regulators could require trust authorities to publish consumers' complaints about certified sites, or regulators could receive and tabulate such complaints. (Analogously, the Department of Transportation tracks and summarizes consumer complaints about airlines.) The resulting transparency would help assure that trust authorities appropriately investigate problems brought to their attention.

For those favor who prefer self-regulation over direct government intervention, BBBOnLine's Privacy seal offers a possible way forward, boasting members more trustworthy than average sites. BBB's tradition of self-regulation seems to help – creating institutional protection against lax review, and blunting short-run incentives to issue unearned certifications. BBB also benefits from its regional network of evaluators, whose proximity to applicants lets them better assess trustworthiness. Yet BBB's small member list and apparent delays make it an unlikely solution to the full problem of online safety. Indeed, after the completion of a draft of this article, BBB closed the Privacy program to new applicants – seemingly a response to the limited number of sites that had chosen to participate in that program. BBB's separate Reliability seal features far more members (some fifty thousand), but with correspondingly less scrutiny on each member. In separate analysis, I found that BBB's Reliability members are somewhat less trustworthy than its Privacy members – providing further reason to doubt whether the BBB's Privacy approach can scale to certify dramatically more sites.

My analysis offers practical lessons for regulators, users, and trust authorities. Regulators should not assume that self-regulatory bodies will assess would-be members correctly or that online publishers will exclude deceptive ads, for incentives are distorted by the profits resulting from selling certifications and advertising. Meanwhile, users should also be wary of supposed certifications – questioning why sites boast of certification, and never assuming that a site is trustworthy merely because the site has obtained an ad slot or even a certification. Finally, trust authorities and advertising platforms might rightly reconsider their practices – realizing that, in the long run, users will come to disbelieve certifications and offers that are tainted by untrustworthy participants.

Acknowledgements

I thank seminar participants at Harvard University's Department of Economics, Business School, and Department of Computer Science, and at the 2006 Workshop on the Economics of Information Security (University of Cambridge). I am grateful to Robert Akerlof, Ross Anderson, Peter Coles, Chris Dixon, Andrei Hagiu, Ariel Pakes, David Parkes, Al Roth, Stuart Schechter, and anonymous reviewers for helpful comments and suggestions.

Appendix A

A.1. Reversibility of conditionals in Bayes rule analysis, when signal and outcome 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 $a + b + c + d = 1$.

	s	$-s$
t	a	b
$-t$	c	d

The definition of conditional probability yields the following identities:

$$P(s|t) = \frac{a}{a+b} \tag{A.1}$$

$$P(s|-t) = \frac{c}{c+d} \tag{A.2}$$

$$P(t|s) = \frac{a}{a+c} \tag{A.3}$$

$$P(t|-s) = \frac{b}{b+d} \tag{A.4}$$

Suppose $P(s|t) < P(s|\bar{t})$. Substituting from (A.1) and (A.2), then cross-multiplying and expanding:

$$\frac{a}{a+b} < \frac{c}{c+d} \tag{A.5}$$

$$ac + ad < ac + bc \tag{A.6}$$

Subtracting ac from each side, adding ab to each side, and regrouping:

$$ab + ad < ab + bc \tag{A.7}$$

$$\frac{a}{a+c} < \frac{b}{b+d} \tag{A.8}$$

Substituting, using (A.3) and (A.4):

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

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

References

Animesh, A., Ramachandran, V. Quality uncertainty and the performance of online sponsored search markets: an empirical investigation. *Information Systems Research*, 21, 1, 2010, 190–201.

Akerlof, G. The market for 'Lemons': quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84, 4, 1970, 488–500.

Backhouse, James et al. A question of trust – an economic perspective on quality standards in the certification services market. *Communications of the ACM*, 2005.

Please cite this article in press as: Edelman, B. Adverse selection in online "trust" certifications and search results. *Electron. Comm. Res. Appl.* (2010), doi:10.1016/j.elerap.2010.06.001

Backhouse, James. et al. Spotting lemons in the PKI market: Engendering trust by signaling quality. In Michael Shaw and M.E. Sharpe Inc. (eds.), *Electronic Commerce and the Digital Economy*. *Advances in Management Information Systems Series Editor*, Vladimir Zwass, 2005.

Baye, M., and Morgan, J. Red queen pricing effects in e-retail markets. SSRN working paper 655448, 2003. Last accessed on May 08, 2009.

Boutin, P. Just how trusty is TRUSTe? *Wired* (April 9, 2002). Available at http://www.wired.com/news/exec/0,51624-0.html, 2002. Last accessed on May 8, 2009.

Brin, S., and Page, L. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, 30, 1–7, 1998, 107–117.

Chiappori, P.-A., and Salanie, B. Testing for asymmetric information in insurance markets. *Journal of Political Economy*, 108, 2000, 56–78.

Claburn, T. Search engines accused of being soft on scammers. *Information Week*. Available at http://www.informationweek.com/news/internet/search/showArticle.jhtml?articleID=193105731, 2006. Last accessed on May 19, 2010.

Consumer Reports Web Watch. A matter of trust: what users want from Web sites. Available at http://www.consumerwebwatch.org/pdfs/a-matter-of-trust.pdf, 2002. Last accessed on May 8, 2009.

Electronic Frontier Foundation. *Letter to the FTC*. Available 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, 1999. Last accessed on May 8, 2009.

Finkelstein, A., and Poterba, J. Adverse selection in insurance markets: policyholder evidence from the U.K. annuity market. *Journal of Political Economy*, 112, 1, 2004, 183–208.

FTC. Self-regulation and privacy online. Available at http://www.ftc.gov/os/1999/07/privacy99.pdf, 1999. Last accessed on May 8, 2009.

Gaudeul, A. Internet intermediaries' editorial content quality. Economics working paper archive, WUSTL Industrial Organization, 2004.

Genesove, D. Adverse selection in the wholesale used car market. *Journal of Political Economy*, 101, 4, 1993, 644–665.

Goodard v. Google. Case No. 5:2008cv02738. California Northern District Court.

Google. Advertising content policy. Available at http://adwords.google.com/support/aw/bin/static.py?page=guidelines.cs. Last accessed on May 19, 2010.

Google. Basics. Available at http://www.google.com/support/webmasters/bin/answer.py?answer=70897. Last accessed on May 19, 2010.

Greenstadt, R., and Smith, M. Protecting personal information: Obstacles and directions. In *Proceedings of the Fourth Annual Workshop on Economics and Information Security*, Cambridge, MA, 2005.

Hansell, S. Paid placement is catching on in web searches. *New York Times*, June 4, 2001.

Jamal, K., Maier, M., and Sunder, S. Privacy in E-commerce: developing of reporting standards, disclosure, and assurance services in an unregulated market. *Journal of Accounting Research*, 41, 2, 2003, 285–309.

Joachims, T., Granka, L., Pan, B., Hembrooke, H., and Gay, G. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the Conference on Research and Development in Information Retrieval*, 2005.

LaRose, R., and Rifon, N. Your privacy is assured – of being invaded: web sites with and without privacy seals. Available at http://www.msu.edu/~larose/es2003post.htm, 2002. Last accessed on May 8, 2009.

Lizzeri, A. Information revelation and certification intermediaries. *The RAND Journal of Economics*, 30, 2, 1999, 214–231.

Marable, L. Consumer reaction to learning the truth about how search engines work. Available at http://www.consumerwebwatch.org/pdfs/false-oracles.pdf, 2003. Last accessed on May 8, 2009.

Microsoft. Online privacy notice highlights. Available at http://privacy.microsoft.com/. Last accessed on May 8, 2009.

Ryan, S. 'Free IPod' takes privacy toll. *Wired*. March 16, 2006.

Tang, Z., Hu, Y., and Smith, M. Protecting online privacy: Self-regulation, mandatory standards, or caveat emptor. In *Proceedings of the Fourth Annual Workshop on Economics and Information Security*, Cambridge, MA, 2005.

TRUSTe. TRUSTe case study – realty tracker. Available at http://www.truste.org/pdf/Realty_Tracker_Case_Study.pdf. Last accessed on May 8, 2009.

TRUSTe Fact Sheet. Available at http://www.truste.org/about/fact_sheet.php. Last accessed on May 8, 2009.

TRUSTe Program Requirements. Available at http://www.truste.org/requirements.php. Last accessed on May 8, 2009.

TRUSTe Watchdog Reports. Available at https://www.truste.org/consumers/watchdog_reports.php. Last accessed on May 8, 2009.

Villeneuve, B. Mandatory pensions and the intensity of adverse selection in life insurance markets. *The Journal of Risk and Insurance*, 70, 3, 2003, 527–548.

Webmasterbrain Search Engine Experiment. Available at: http://www. webmaster-brain.com/seo-tools/seo-experiments/the-search-engine-experiment/test-results/, 2006. Last accessed on August 20, 2007.

Please cite this article in press as: Edelman, B. Adverse selection in online "trust" certifications and search results. *Electron. Comm. Res. Appl.* (2010), doi:10.1016/j.elerap.2010.06.001