Who's Naughty? Who's Nice? Social Preferences in Online Industries

Mitchell Hoffman and John Morgan
UC Berkeley; UC Berkeley and Yahoo *

February 2011

Abstract

We conduct a battery of social preference experiments on business people from two controversial internet industries—domain trading and adult entertainment (pornography). We then conduct the same experiments on Berkeley students. Levitt and List (2007) conjecture that selection pressures among business people would reduce or eliminate pro-social choices. We find the opposite: Internet business people were significantly more altruistic, more trusting, reciprocate more, lie less, and respond differently to shame than students. We present a model of reverse selection that rationalizes our findings.

*We thank Stefano Della Vigna, Shachar Kariv, and Michael Schwarz for helpful suggestions. We are especially grateful to Steve Tadelis for detailed comments and advice. We are indebted to the conference organizers Jay Kopita, Aaron Kvitek, and Connor Young for allowing us to conduct the experiments. Hoffman gratefully acknowledges support from the National Science Foundation IGERT Fellowship. Morgan gratefully acknowledges support from the National Science Foundation through grant NES-0095639. Address correspondence to: Morgan (morgan@haas.berkeley.edu): Haas School of Business, Berkeley, CA 94720-1900.

1 Introduction

The prevalence of pro-social preferences in influencing choice behavior is one of the most important findings to emerge from laboratory experiments. Laboratory subjects are more generous, more trusting, more trustworthy, and more truthful than models where decisions are guided purely by pecuniary considerations would suggest. A key question for businesses and policymakers is whether these results are relevant to domains outside the lab and to individuals other than undergraduate students. That is, to what extent are these findings externally valid?

This matter has been hotly debated. The arguments against external validity emphasize the artificiality of the laboratory environment, the non-representative sample, and the relatively small stakes of the experiment, as well as selection pressures in real-world populations. These critiques are nicely summarized in Levitt and List (2007), who write (pp. 168-169):

> *In contrast to the lab, many real-world markets operate in ways that make pro-social behavior much less likely. In financial markets, for instance, the stakes are large, actors are highly anonymous, and little concern seems to exist about future analysis of one's behavior. Individuals with strong social preferences are likely to self-select away from these markets, instead hiring agents who lack such preferences to handle their financial dealings. Thus, one must take great care when claiming that patterns measured in the experimental economics laboratory are shared broadly by agents in certain real-world markets.*

Arguments supporting lab-based evidence on social preferences emphasize the robustness of pro-social behavior across differing subject pools (e.g. Fehr and List (2004); Roth et al. (1991)), with high stakes (e.g. Slonim and Roth (1998)), and in natural contexts (e.g. Karlan (2005); Schechter (2007)).[1] Several studies have demonstrated that for a given subject population, other regarding behavior in the lab correlates with the same behavior in the field. For instance, in a study of charitable giving by students, Benz and Meier (2008, p. 270) write:

> *Overall, however, our results lend remarkable support to the notion that behavior in an artificial experiment corresponds to students' behavior outside the laboratory. ..this suggests that experiments can provide useful information about behavior in the field.*

Our paper contributes to this debate. We report the results of experiments comparing choice behavior between students and internet business people. Our experiments shed light on many aspects of social preferences including altruism, trust, reciprocity, lying, and shame. While altruism, trust, and reciprocity have been widely studied, lying and shame have received less attention. To

[1] Camerer (2003) provides a comprehensive study of experiments on other regarding preferences including field experiments. More recently, Cardenas and Carpenter (2008) provide a useful review of pro-social behavior in lab experiments conducted in developing countries.

the best of our knowledge, we are the first to conduct field experiments on these aspects of social preferences.

Perhaps more importantly, we contribute to the debate by examining the selection hypothesis—the notion that individuals exhibiting social preferences select out of cutthroat settings. To examine this hypothesis, we chose subjects from two industries where we would expect such selection would operate forcefully—internet domain trading and online adult entertainment. Barriers to entry and exit are low in both of these industries, making them highly competitive. Both industries gained notoriety in recent years for various unethical practices including cybersquatting and making pornographic material easily available to minors. Thus, we would think they would be fairly unattractive to individuals with pro-social preferences. We compared choice behavior of subjects in these two industries with Berkeley undergraduate students participating in identical experiments in a standard laboratory setting.

We find significant differences in the choice behavior of internet business people and students in the lab—but not in the expected direction. Internet business people are more altruistic, more trusting, and lie less. They also respond differently to extrinsic incentives for good behavior (i.e. monitoring and shame).

What accounts for these differences? We argue that there is another form of selection in the real world—especially in online settings. Since solving the trust problem with formal mechanisms is difficult, only those individuals perceived to be trusting and trustworthy can avail themselves of informal mechanisms and ultimately succeed. Thus, our data is consistent with precisely the opposite selection compared to the conventional wisdom.

Comparing internet business people to students in the lab varied two components: the subject populations and the setting in which the experiments were conducted. Our experiments on internet business people were conducted in the field at industry conferences outside of laboratory settings. To examine differences between behavior in the field and in the lab, we subsequently ran an additional set of experiments on Berkeley undergraduates, this time conducted on a popular student thoroughfare. The difference in setting matters: experiments conducted with students in the field produced more evidence of pro-social preferences than these same experiments conducted in the lab. Compared to field experiments with students, internet business people still tend to be more pro-social, but the differences are reduced compared to the behavior of students in the lab. Taken together, our main findings suggest that: (1) placing students in artificial laboratory settings tends to *understate* the degree of pro-social preferences; and (2) internet business people tend to be more pro-social than students. Neither finding is consistent with the view that pro-social preferences observed in the laboratory overstate the degree of pro-sociality "in the wild."

The remainder of the paper proceeds as follows. Section 2 provides background about the online adult entertainment and domain name industries. Section 3 describes our experimental design. Section 4 gives the results of our experiments. Section 5 discusses several explanations for

our results. Section 6 concludes.

2 Industry Background

The popular press has highlighted a number of unsavory practices online. Most notably, the Internet is widely perceived to have made sexually explicit materials considerably more accessible—especially to minors. These concerns were enough to motivate Congress to pass several laws including, most recently, the Children's Internet Protection Act. Another widely criticized practice is "cybersquatting"—the registration of domain names with the intent of profiting from an unrelated firm's brand or goodwill. For instance, the domain "wwalgreens.com" clearly seeks to profit from association with the brand and goodwill of Walgreens though its owner is in no way related to the drugstore.

Who enters these types of businesses? Online adult entrepreneurs would not seem especially concerned about negative externalities from others' exposure to their materials. Likewise, the motives of a domain trader who profits from other firms' brands and trademarks hardly seem trustworthy or altruistic.

Before proceeding further, it is useful to provide a more detailed overview of these two industries.

2.1 Online Adult Entertainment

Adult entertainment is big business on the internet. Online adult entertainment produced annual revenue of $2.8 billion in 2008, which accounts for about 21% of all adult entertainment revenues (Edelman, 2009). Ropelato (2006) reports that 12% of all websites are pornographic. In addition to website owners and entrepreneurs, there are website designers, marketers, models/actors, as well as back office workers.

Industry revenue has been generated primarily through subscriptions to adult websites providing sexual videos and images. Subscription sites have been recently challenged by adult video-sharing sites where adult videos can be downloaded for free (Edelman, 2009). Sexual themed social networking websites, such as AdultFriendFinder.com, and websites specializing in live video camera content ("cams") represent new and fast-growing categories. Indeed, AdultFriendFinder.com, was ranked in July 2009, according to Alexa.com, as the 71st ranked most popular website on the internet, and was sold in 2007 for $500 million (Hopkins, 2007).

Generating leads is a key part of the business. While the consumer only pays at the subscription website, the firms that directed the consumer to the final website often receive significant bonuses, frequently equal to several times the monthly subscription rate.

This industry has been the subject of fierce public criticism. It is argued that online pornography is offensive to women, lowers men's perceptions of women's abilities, is addictive and time-consuming, negatively affects relationships by creating unreasonable sexual expectations, distorts

views of normal sexual practices, and exposes children to inappropriate sexual material. Two especially serious criticisms are leveled at the industry: that it promotes sexual assault and rape, and that minors are among the performers.[2] Many of the above criticisms are substantiated by scientific research.

2.2 Domain Name Industry

The domain name industry consists of businesses involved in buying, selling, and developing internet domain names. In 2007, it was estimated there were over 10,000 domain investors worldwide (Kesmodel, 2008, p. 11).

Revenues in this industry derive from two main sources: advertising and sales of domains to third parties. Firms earn revenues through pay-per-click (PPC) advertising. Specifically, an advertiser places its ad on the domain and pays the owner for each time the ad is clicked. Most of the traffic comes from "direct navigation," where consumers type the web address directly into their internet browser's URL bar. For example, a consumer looking to purchase a vacation in Florida may type "FloridaVacation.com" into her URL bar instead of using a search engine. Revenues also stem from domain resale. While infrequent, domain resale can be incredibly lucrative, with names often selling for many times the purchase price.[3] In 2007, the size of the industry was estimated at $2 billion annually (Goldman, 2007).

The domain industry has also seen its share of controversy. As mentioned above, some domain traders have registered domain names which include company trademarks or typos of trademarks, a practice known as "cybersquatting" or "typo-squatting". This was particularly problematic in the mid 1990s, and led to regulation in the form of the Anticybersquatting Consumer Protection Act of 1999 (Kesmodel, 2008). Despite this legislation, cybersquatting still does occur, and, may even be on the rise.[4].

[2]For example, on the subject, Zillmann and Bryant (1992) show that large-scale lab exposure increases callousness toward rape victims. However, using cross-state data, Kendall (2006) argues that online pornography may have decreased rape, though the study uses data on total internet usage instead of usage of online pornography, and there is limited state-wide variation in the timing of internet penetration. On the disturbingly large problem of child pornography on the internet, see Jenkins (2001).

[3]In unscientific interviews conducted in 2006 and 2007, most domain traders reported selling between 1 and 5 percent of domains in their portfolios every year.

[4]For example, in July 2009, comedian Jay Leno won the domain "thejaylenoshow.com" in a domain arbitration case from a person named Guadalupe Zambrano who had registered the name first in 2004. See also several comments by Benjamin Edelman on this topic (Edelman, Benjamin, 2003a,b; Edelman, 2009). Domain traders also profit when users forget to re-register valuable domain names. If the current owner of a domain forgets to re-register, and a short grace period elapses, the name goes up for auction or becomes available for public registration. In some cases, traders may attempt to sell the name back to the original owner at a large mark-up. Domain traders vigilantly monitor this market. For example, in the early 2000s computer-savvy domain traders such as Yun Ye scooped up scores of names that unsophisticated owners forgot to re-register and profited handsomely (Kesmodel, 2008)

3 Experimental Design

We conducted a series of experiments designed to probe three key aspects of other-regarding behavior: altruism, shame-aversion, and lying aversion. All of these aspects have featured prominently in earlier work on undergraduate populations.

Game 1: Trust with Shame

In this game, Player A decides whether to choose IN (cooperate) or OUT (defect). If A chooses OUT, the game ends and each player receives \$5. If A chooses IN and B chooses DON'T ROLL (defect), A receives \$0 and B receives \$14. If A chooses IN and B chooses ROLL (cooperate), B receives \$10 and the payoff to A is decided by chance. With probability 5/6, A receives \$12, and with probability 1/6, A receives \$0.

In the baseline treatment, Player A is merely informed of his or her payoff. In the shame treatment, Player A is informed of his or her payoff as well as the action selected by Player B. The baseline version of game 1 is a binary version of a standard trust game.[5] The shame treatment appears in a recent paper by Tadelis (2009).

It is easy to see that, regardless of the treatment, the unique subgame perfect equilibrium for self-interested players is for A to choose OUT and for B to choose DON'T ROLL. A player with altruistic preferences is more likely to be trustworthy (i.e. choose ROLL) and to be trusting (i.e. choose IN). With shame averse players, trustworthy and trusting behavior should both increase.

We conducted this experiment in early 2008 at the DOMAINfest Global domain name conference—a trade show for domain investors and others in the industry. The conference organizers allowed us to set up a booth in the exhibition hall. To attract potential subjects, we hung a large sign saying "Economics Experiment" on the front of the booth and displayed a briefcase filled with money on the table. The experiment was conducted in paper form using the strategy method. Subjects read the instructions and then made their decisions both as Player A and as Player B (strategy-method). Afterwards, subjects were given an additional sheet asking for demographic information and belief elicitation. Fifty-two subjects participated in the experiment.

Subjects were instructed to return after a specified time in the afternoon to pick up payment. In the interim, subjects were randomly matched and assigned a random role. We then computed each subject's payoffs and paid them privately in cash and in person at the end of the conference day. Subjects in the shame treatment who received \$0 as Player A were informed at the time of payment whether or not Player B chose ROLL. On average, subjects earned \$15.97 for the experiment.[6] The experiment form took about 5 minutes for a typical subject to complete, amounting to an hourly payment rate of \$191.67 per hour.

It is worth making a general comment about the payment and timing of our experiments. The

[5] The trust game was first studied in Joyce et al. (1995).

[6] The payment amount was determined in conjunction with the conference organizers and included a \$7 show-up fee.

most important issue for us in designing payment for all our experiment was that the stakes would be salient for internet business people. To participate in the conference, we were asked to conduct a short experiment emphasizing that participants would receive cash. To make payments salient while preserving reasonable stakes, we eliminated virtually all wait time from the internet business people experiments, allowing subjects to make decisions almost immediately after arriving at our table. Thus, internet business people earned very high earnings per hour, even while their total earnings were fairly similar to that of students.

We repeated the same experiment using an identical number of subjects at the Berkeley Experiment Social Science Laboratory (Xlab). Here, the subject pool consisted of undergraduate and graduate students enrolled at UC Berkeley. Once again, subjects filled out forms playing both roles. Unlike the experiments conducted at the domain conference, where subjects showed up individually over the course of the day, here all subjects participated at the same time. This design resulted in considerably more wait time for subjects—it took approximately one hour to complete. Subjects earned an average of $12.77.

Game 2: Altruism and Lying

In this game, Player 1 decides between two messages to send to Player 2: "Option A will earn you more money than Option B" (Message 1) or "Option B will earn you more money than Option A" (Message 2). If Option A is selected, Player 1 earns less money and Player 2 more than if Option B is selected. Thus, Message 1 is true (but costly to Player 1 if acted upon) while Message 2 is false. After receiving the message, Player 2 chooses between the two options. Player 2, however, has no information about the game other than the message.

We used three different payoff treatments for the game: ($5, $6) versus ($6, $5), ($5, $6) versus ($15, $5), and ($5, $15) versus ($15, $5). The first number denotes Player 1's payoff while the second denotes Player 2's.

If a self-interested Player 1 believes Player 2 will follow his advice, then his best strategy is to send Message 2 in every case. More generally, a self-interested Player 1 will choose whichever message is more likely to lead Player 2 to choose Option A.

As a comparison, we also employ a simple binary dictator game using the same payoffs. Here, a self-interested Player 1 will simply select Option B in every case.

Gneezy (2005) proposed this pair of games as a test of altruism versus lying aversion. In the message game, Player 1 might select the "altruistic" Message 1 either out of concern for the financial payoffs of Player 2 or because of negative payoffs from having to mislead Player 2 to obtain the preferred Option B. The dictator treatment is a pure test of altruism—the only reason to choose Option A is out of concern for Player 2's payoffs. Thus, within-subject differences in behavior may represent lying aversion.

We performed the experiment at Cybernet Expo in San Francisco in late spring 2008. Cybernet Expo is a medium-sized trade show for people in the online adult entertainment industry. The

conference organizers provided us with a table near the registration desk. Our setup was similar to that used at DOMAINFest. Forty-four subjects participated. The experiment lasted around 5-10 minutes and subjects received an average of $15.76.[7]

Since our focus in this experiment is on the actions of the Sender (where we observe trade-offs between lying aversion and money), we matched four senders to one receiver, and four dictators to one recipient. Each Player 1 made six total decisions: three message-sending decisions and three dictator decisions. The subject was informed that his message or dictator decision would be matched with another person without reference to the possibility of multiple matches to the same Player 2. This helps increase our sample size of subjects in the Player 1 role without altering the competitive dynamics of the game.

We then repeated this experiment with undergraduate subjects in the Xlab. We conducted these experiments with 74 students divided into two large groups, instead of one person at a time as at the conference. This increased the duration of the game and, consequently, lowered the payoffs per unit time. Average earnings were $12.93 for a session that lasted around 1 hour.

Game #3: Altruism, Shame, and Lying: The Envelope Game

In this game, each subject is given an unmarked envelope containing 5, 10, 15, 20, 25, or 30 one-dollar bills, with each number equally likely. Subjects then remove however many bills they wish to keep. Subjects then seal the envelope containing the remaining bills (if any) and write a number on the outside to indicate the number of bills originally inside. The exact wording was:

> *Next, you will send a message to whoever receives your envelope later today by writing a number on the outside of the envelope. Your message will indicate the number of dollar bills originally inside the envelope; however, the other person will not be informed of the actual number of dollar bills originally contained in the envelope. They will only be informed of the number you write down.*

Subjects then place their envelopes in a box, where they are randomly shuffled. Finally, each subject draws an envelope from the box and keeps the money.

Three different margins are at play in the envelope game: altruism, shame aversion, and lying aversion. Since we are examining multiple margins, one needs a bit more structure on preferences and, in particular, on interaction effects between these margins to derive predictions. Motivating our hypotheses is the following, fairly general, model of preferences: Suppose that utility is of the form:

$$U = U(E - \pi, \pi, E - M, M - \pi; L, S, \theta).$$

where E is the subject's endowment, π is the amount they give to the other player, and M is the amount they write on the envelope to indicate their initial endowment. The first argument

[7] For payment, we use the original payment schedule of Gneezy (2005) plus a $5 show-up fee. The amount paid is equivalent to between $94.56 and $189.12 per hour.

is payoff to oneself, the second argument is payoff to the other person, the third argument is the magnitude of a person's lie, and the fourth argument is the difference between one's endowment and what one gives (over which a person may incur shame). L is a parameter for lying aversion, S is a parameter for shame aversion, and θ is a parameter for social preferences. We assume that utility is increasing in the first two arguments, decreasing in the latter two, and strictly concave in each of its arguments. Finally, we assume that an increase in θ lowers the marginal utility of own payoff and raises the marginal utility of the payoff to the other person. Similarly, an increase in L increases the marginal disutility of a given lie $(E - M)$ while an increase in S raises the marginal disutility of a given magnitude of shame $(M - \pi)$. This model offers the following predictions about behavior:

Consider an individual who optimizes at an interior solution. Such an individual optimally gives some fraction of her endowment and reports an endowment amount below the truth. When that individual becomes more altruistic, it has the direct effect of leading her to increase giving. The indirect effect is to reduce the difference between the reported endowment and the true endowment. The reason is that the increased giving has reduced the marginal loss associated with shame. Equating the marginal cost of shame and lying requires more truthful reporting. Thus, increasing altruism increases both giving and truthfulness.

Using similar logic, an increase in lying aversion has the direct effect of increasing truthfulness and the indirect effect of increasing giving. An increase in shame aversion, on the other hand, increases giving but reduces truthfulness. To the best of our knowledge, this is the first experiment incorporating all three effects into a single design.

We also conducted this experiment at Cybernet Expo. Throughout the day, subjects participated in the first part of the experiment—removing cash and then returning a sealed envelope to the box. Subjects were told to return at the end of the day for the second part of the experiment—drawing an envelope from the box. It took subjects about 5 minutes to complete the experiment. The average earnings for the 54 subjects who participated in the experiment were \$18.98, which works out to an hourly rate of \$227.76.

We repeated the experiment with 54 undergraduate subjects at Xlab. As in the other experiments, we used the same payment structure, but reduced payment per hour by conducting the experiment in a large group instead of one at a time. Privacy is obviously a key concern in dictator-like games. Cybernet subjects made their decisions alone at a separate table. To mimic this in Xlab, subjects exited the room one at a time to make their decision and then returned subsequently. Average earnings were \$18.15 and sessions lasted between 45 minutes and one hour.

4 Results

Though we performed different experiments on domain traders and adult entertainment professionals, the observed behavior, compared with students, is very similar. In virtually all cases, internet

business people exhibited higher levels of other-regarding behavior.

Our main findings are about internet business people compared to students are as follows:

R1. Internet business people behave more altruistically than students.

R2. Internet business people are more trusting than students.

R3. Internet business people are less likely to lie than students.

R4. Internet business people do not appear affected by shame. In fact, for internet business people, shame may even crowd out pro-social behavior.

We now present the evidence for these findings.

R1. Internet business people behave more altruistically than students.

The results from Games 1-3 all provide supportive evidence.[8]

In the dictator treatment of game 2, students chose the selfish option 83% (90/108) of the time compared to 61% (68/111) of the time for adult industry professionals. This difference is highly significant according to a two-sided Mann-Whitney test ($z = 3.64$, $p < 0.01$), which we will use as our primary test of significance.[9] Across the three payoff treatments, the percentages choosing the selfish option were (70%, 32%, and 81%) for adult professionals and (97%, 56%, 97%) for students. The difference across each payoff treatment is also highly significant ($p < 0.01$, $p = 0.05$, $p = 0.03$).

In the envelope game (Game 3), students gave an average of 14% of their endowments, far less than adult professionals, who gave 39% on average. This difference is highly significant ($z = 4.23$, $p < 0.01$). Figure 1 below displays giving by endowment level. As the figure shows, giving levels are higher for adult professionals than students at every level. The differences are most pronounced when endowments are extremely high or extremely low.

The second stage of the trust game (Game 1) may also be used to measure altruistic behavior. In this stage, students chose ROLL (the trustworthy/altruistic option) 33% of the time, far less frequently than domain traders who chose this option 52% of the time. Once again, the difference is significant ($z = 1.98$, $p < 0.05$).

Our adult professional population is predominantly male while our student population is more evenly distributed by gender. A number of studies have reported gender differences in altruistic behavior (see, e.g., Eckel and Grossman (1998), as well as Andreoni and Vesterlund (2001)). Their main finding is that women behave more altruistically than men. Thus, it might be important to control for these gender differences across the two populations.

Table 2 below presents linear probability models of altruistic behavior controlling both for the subject population and gender. The first column of the table compares altruistic behavior in game 3 and shows that adult professionals give 25 percentage points (almost 200%) more of their

[8] We will say that someone is behaving altruistically when they sacrifice their own monetary payoff to increase the monetary payoff of the other player without any possibility of subsequent positive monetary return. Thus, trusting in a trust game is not (necessarily) altruistic.

[9] While the text reports only Mann-Whitney test statistics and p-values, we also performed hypothesis testing using t-tests assuming unequal variance and Fisher exact tests. Our conclusions are unaffected by the test statistic employed.

endowments than do students. In contrast, gender has no effect on giving.[10] The second column of the table compares trustworthy behavior in game 1. Once again, the effect of adult professional status is large and highly significant. Here, we also find a significant gender effect—men are considerably less likely to be trustworthy.

R2. Internet business people are more trusting than students.

Examining trust in Game 1, we find that 67% of domain traders chose IN, the trusting option, compared to 38% of students. The difference is significant at the 1% level ($z = 2.93$).

Two factors influence the decision to be trusting. First, if one believes that trust is likely to be reciprocated, then the trust option is payoff maximizing. The second factor is altruism. Even if one believes that trust will not be reciprocated, then choosing the trusting option is, in effect, making a payoff reducing transfer to the other player.

About 52% of domain traders believe in the trustworthiness of the other person compared with only 31% of students. Thus, the higher levels of trust displayed by domain traders might simply represent a payoff maximizing response to these beliefs. However, 48% (10/21) of domain traders chose the trusting option even when they thought the other person was likely to be untrustworthy. Only 25% (9/36) of students chose the trusting option under these same beliefs. The difference is significant at the 10% level ($z = 1.73$, $p = 0.08$). Thus, the higher levels of trust shown by domain traders are not simply a reflection of payoff maximizing behavior.

Returning to Table 2, we compare trust in the two populations adding controls for gender and beliefs. As column 3 of that table shows, trust is correlated with positive beliefs about trustworthiness. However, even with this control, the level of trust shown by domain traders is 24 percentage points higher than that by students. Once again, gender has no effect on trust.

R3. Internet business people are less likely to lie than students.

In the cheap talk game (Game 2), a sender may wish to "lie" to the receiver by offering "wrong" advice (i.e. message 2). If lying is defined solely as choosing Message 2, adult industry professionals and students have similar lying rates (55/111 vs. 47/111; $p = 0.28$).

This, however, ignores the sender's intent in sending message 2. If the sender believes that his advice will be taken, then message 2 is intended to deceive the receiver into choosing the lower payoff action. If, however, the sender believes that the opposite of his advice will be taken, message 2 constitutes a truthful message since it induces the receiver to choose the higher payoff action. In that sense, message 1 constitutes a lie. Our preferred definition of lying takes into account the beliefs that the message will be followed.

In Game 2, we elicited these beliefs. If the sender believes that the receiver will follow his advice, we code message 2 as a lie. If the sender believes the opposite, we code message 1 as a lie. With this more nuanced version of lying, the results are vastly different. Adult industry professionals lie only 46% (36/78) of the time compared to 71% (53/75) of the time for students ($z = 3.06$, $p < 0.01$).

[10] Interacting gender with subject population also produces an insignificant coefficient.

One can also define lying more conservatively by throwing out all subjects who didn't believe that others would follow their advice. Under this definition, adult professionals also lie significantly less than students (30/66 for adults vs. 21/30 for students, $z = 2.22$, $p = 0.03$).

In Game 3, subjects can lie by writing a number on the envelope different from their original endowment. They might be motivated to do this to appear more generous in their giving (i.e. to lessen the shame of appearing to be ungenerous).

Since lying is only possible for endowments in excess of the minimum ($5), we restrict attention to choices made when lying was feasible. Here we find that students lie 33% of the time compared with a 22% lying rate for adult professionals. Despite the fact that students lie 50% more often than adult professionals, this difference is only significant at the 24% level ($z = 1.17$).

To summarize, adult professionals lie slightly less when the tradeoff is between lying and shame, and significantly less when the tradeoff is between lying and money.

R4. Internet business people do not appear affected by shame. In fact, for internet business people, shame may even crowd out pro-social behavior.

The main treatment in Game 1 was whether the choice of the second player was disclosed or not. Disclosure adds the possibility of extrinsic motivation (shame) for acting in a trustworthy fashion. Table 3 shows the percentage of trusting and trustworthy behavior under the two treatments for each population.

The effects of the shame treatment produce opposite results in students and domain traders. Domain traders are half as likely to be trustworthy under the shame treatment—a highly significant difference ($z = 2.18$, $p = 0.029$)—and one that goes in the opposite direction from our priors. Domain traders appear to anticipate this effect and factor it into their decisions to trust in the first place. The amount of trusting behavior falls by about one-sixth—however the difference is not statistically significant ($p = 0.38$).

Among students, the effect goes in the anticipated direction, but is weak. Both trust and trustworthiness modestly increase under the shame treatment; however the differences are insignificant ($p = 0.77$ for trustworthiness and p = 0.57 for trust).[11]

Why does the shame treatment reduce pro-social behavior among internet business people? Clearly this population is not strongly driven by shame aversion. One possible explanation is that extrinsic incentives are crowding out intrinsic incentives for altruism. That is, we might be observing an effect similar to that demonstrated in Gneezy and Rustichini (2000) as well as the recent paper by Ariely et al. (2009) where the introduction of cash payments can actually decrease pro-social behavior by crowding out altruistic motivations. In our case, a small extrinsic punishment may crowd out intrinsic altruistic motivations.

[11] The absence of a significant effect for students may be driven by the strength of the treatment. Tadelis (2009) finds evidence that shame can induce pro-social behavior; however the magnitude of his effects varies with the treatment strength. We use the same control as Tadelis and his mildest shame treatment as our treatment and observe effects of similar magnitudes. In our study, this does not rise to the level of statistical significance while in his study it did.

Replication

The results above suggest that internet business people exhibit more other regarding behavior than Berkeley students. One alternative explanation for the difference in behavior is that our population of Berkeley students exhibits dramatically less other regarding behavior than typical student populations.

As stated above, earlier researchers reported results from games 1 and 2. Thus, to examine this alternative explanation, we compare other regarding behavior in our student experiments with earlier results. Tadelis used Berkeley students in his experiments, which correspond to game 1. Comparing trusting and trustworthy behavior in the two sets of experiments, we fail to reject the null hypothesis that we can pool the two student populations at conventional levels.[12]

Gneezy used Israeli students in his experiments, which correspond to our game 2. For the message game treatments, we fail to reject the null hypothesis that the samples may be pooled. In the dictator treatments, we fail to reject the null of pooling for two of the three sets of parameter values.[13]

Overall, these results suggest that we were able to successfully replicate earlier experiments with student populations. Hence, differences in behavior between internet professionals and students are not merely an artifact of unusually selfish choice behavior on the part of Berkeley students.

5 Discussion

Why do internet business people exhibit more other regarding behavior than students? Common wisdom suggests that pro-social characteristics would be selected away from in the anonymous online world. However, it is precisely in semi-anonymous situations where trust is at a premium—suggesting the opposite selection.

For example, domain traders frequently engage in deals supported by limited formal contracting instruments. A very common way to transact a domain name is for the buyer to send cash to the seller, and then for the seller to transfer the name. An unscrupulous seller, though, could decide simply not to transfer the name after receiving the money. While such problems could occur in many industries, they present more of a potential threat to the domain industry. First, domain property law is relatively new.[14] Second, unlike in real estate, there are no title companies and no title

[12] Formally, we compare the fractions of trusting and trustworthy behavior in our sample and in Tadelis under each treatment and compute associated Mann-Whitney test statistics. The p-value in all cases exceeds 0.23.

[13] For the dictator treatment where subjects chose between allocations of (6, 5) and (5, 6), we observe less altruistic behavior in Berkeley students.

[14] The case of Sex.com is a famous illustration of the difficulties of early domain name property law. The domain name Sex.com was stolen from Gary Kremen, a domain trader, by a convicted felon named Stephen Michael Cohen. Cohen stole the name by sending a forged fax to the domain registry Network Solutions; Network Solutions transferred the name without contacting Kremen. Kremen's initial legal efforts to hold Network Solutions accountable for transferring the name were frustrated by the determination of a California district court that domain names, unlike real estate, were merely intangible property, for which the tort of conversion did not apply (Kremen v Cohen 2000). This part of the ruling was reversed (Kremen v. Cohen 2003), though, by the decision of an appeal court in 2003,

insurance.[15] Crucial to the enforcement of domain deals is that the major players in the industry are small in number and relatively tight-knit, and thus even though the internet is anonymous, dishonest behavior can come to light in this market. The usual sanction is the exclusion of a guilty individual from future trading opportunities.

To formalize how pro-social selection in semi-anonymous markets might arise, we offer the following stylized model: Consider a continuum of traders that play a modified prisoner's dilemma each period with a randomly matched opponent. The model consists of two periods. The mass of traders is set to unity, and there is no entry, the idea being that those who do not take advantage of business opportunities at critical times (like the opening of the .com space or the introduction of pornography on the internet) may face serious obstacles to entry later on.

There are two types of individuals: selfish individuals, who are a fraction π of the population, and pro-social individuals. The stage game is given as below:

	C	D	DT
C	b, b	$d, a-\theta_i$	$k, 0$
D	$a-\theta_i, d$	c, c	$k, 0$
DT	$0, k$	$0, k$	$0, 0$

where C stands for cooperate, D for defect, and DT for don't trade. We assume that $a > b > c > 0 > d$ and that $k \leq 0$. The parameter k represents a cost associated with unsuccessfully initiating a trade. One can think of this cost as being arbitrarily small. The ordering on the other parameters merely ensures that: (a) the game is a modified prisoner's dilemma; and (b) trading even with mutual defection is preferred to not trading.

The parameter θ reflects an individual's social preferences. Selfish individuals have $\theta_i = 0$ whereas pro-social individuals have $\theta_i = \theta \in (0, a-b)$. The upper bound on θ merely guarantees that pro-social types would not wish to cooperate were this a one-shot game. At the conclusion of the period, payoffs are realized and a fraction $r < 1$ of trades are publicly revealed. That is, trading is semi-anonymous and some defections will escape notice.

In addition, we make the following three assumptions:

Assumption 1 (A1). $rc < a - b < c - d$.

Assumption 2 (A2). $\pi < \min\left[\frac{b-c}{b-d+rc}, \frac{cr}{c-d+cr^2}\right]$.

Assumption 3 (A3). $\theta > \bar{\theta} = \frac{1}{1-\pi}\left(cr(\pi r - 1) + (a-b)(1-\pi) + \pi\,(c-d)\right)$

A1 states that the probability of revelation, r, is sufficiently low, and that the temptation to defect is smaller than the penalty from being defected on. A2 states that the share of selfish players in the population is sufficiently small. A3 states that pro-social players have a sufficiently

where the court took a broader view of domain names as property, eventually leading to a multi-millionaire dollar settlement of Kremen with VeriSign (formerly Network Solutions) in 2004. See McCarthy (2007) for more details.

[15] There are, however, online escrow services that can be used to facilitate a transaction, for example, Escrow.com. It is our impression, though, that escrow services are mostly used for high value transactions.

high degree of pro-social preferences. If there were no or only a few pro-social players, cooperation would be impossible owing to Assumption 1 and continuity.

With these, we are able to prove the following proposition.

Proposition 1 *There exists a Perfect Bayesian Equilibrium (PBE) where pro-social players cooperate and selfish players defect in the first period, and in the second period, players revealed to have defected are not traded with.*

A key empirical implication of this proposition is the following: If we surveyed the population of traders in the second period, we'd find that pro-social types are *overrepresented* in the population. The reason is that opportunistic trading behavior on the part of the selfish types leads (probabilistically) to their exclusion from future trading. Moreover, if all traders were selfish or hired selfish types as their representatives, cooperation in this market would collapse entirely.

Of course, we have merely selected an equilibrium in the game. It is well-known that the modified prisoner's dilemma (including our variant) suffers from equilibrium multiplicity. Is there any reason to expect that the equilibrium described in Proposition 1 is likely to emerge? Our next proposition offers some justification: For a dense set of parameter values, this equilibrium welfare dominates all others. Formally,

Proposition 2 *The PBE identified above is Pareto superior to all other PBE.*

A second form of reverse selection involves the need to work effectively in teams. Undergraduate students have limited opportunities and limited need to work in teams. Indeed, in much of their graded work (problem sets, exams, etc.), working in teams is explicitly forbidden. Thus, individuals lacking pro-social preferences are at no particular disadvantage.

In contrast, teamwork is often essential in industry environments (Lazear, 1998). Success in the adult entertainment industry requires considerable teamwork in the creation, distribution, and marketing of content. Informal alliances also play a key role. For example, internet business people may collaborate on developing a valuable domain name, refer one another to different online advertising companies, or share ideas about optimal pricing / fees. In such team settings, trust and honesty are at a premium. Moreover, among those in online industries, it may be difficult for purely reputational considerations to provide sufficient "glue" to hold together the trust relationship. Individuals with pro-social preferences, however, will find it easier to be good team players.

Quite simply, being branded "not a team player" can spell the end of one's career in corporate settings, but has no such effect in academic settings.

Still another possibility is that the context of the experiments themselves may activate differing preferences. The formal laboratory setting, where students are required to sit apart from one another and not communicate in any way, resembles an exam setting for students. In contrast,

attending an industry conference is inherently a social activity and perhaps merely being in this setting created an atmosphere of pro-sociality.

It has been argued that pro-social preferences seen in laboratory experiments might arise through selection—pro-social students are more likely to volunteer for experiments. Field experiments might exhibit the same or stronger selection pressures. Clearly, misanthropic types are less likely to select to attend industry conferences and volunteer for experiments. For domain traders, however, this selection concerned is allayed somewhat by the fact that DOMAINfest is currently the industry's largest conference, and is attended by many of industry's main players. Selection concerns are greater for Cybernet Expo, which attracts only a small fraction of those in the online adult entertainment industry.

We designed our experiments and selected subject pools based on our priors that internet business people—individuals operating in cutthroat industries—would be unlikely to exhibit strong pro-social preferences. Thus, our paper offers an ideal testbed for the contention in Levitt and List (2007) that pro-social preferences observed in the lab are likely to be selected out in other populations. Having found the opposite result, we offer the above alternative explanations but cannot rule any of these in or out. This remains for future experiments.

6 The Lab versus the Field

The comparisons above suffer from a confound—our experiments with internet business people differ both in the population of subjects (our main treatment) and in the setting (lab versus industry conference). Thus, one possible explanation for our results is that internet business people are, in fact, *less* pro-social than students but that field settings provoke more pro-social responses. While we are unaware of any theory that predicts pro-sociality will be less in the lab than in the field, our initial experimental setup does allow for this objection and several readers raised this possibility.

To separate the two effects, we ran additional experiments in a field setting with a student population.[16] Specifically, we set up a table, much like the ones we used at the industry conferences, at Sproul Plaza, a popular thoroughfare for Berkeley students. We then asked passing students if they wished to participate, much as we did with internet business people. Using this setup, we reproduced games 1 and 2 from our previous design. This yielded a sample of 82 students for Game 1 and 49 students for Game 2. To make the per hour compensation rates the same as our laboratory experiments, we reduced the payouts by a factor of 5 since these experiment took only about one-fifth as long to complete.

Below, we re-examine results 1-4, making students in the field with students in the lab and internet business people.

R5. Students behave more altruistically in the field than in the lab. Internet

[16] The alternative, running laboratory sessions with internet business people, proved unworkable.

business people are as altruistic as students in the field.

First, we analyze student behavior in the lab versus the field in the dictator treatment of game 2. In this game, students in the field chose the selfish option 67% of the time compared to 83% of the time in the lab ($z = 2.93$, $p < 0.01$). In the field, students chose the selfish option 64%, 55%, and 82% across the three payoff treatments compared to 97%, 56%, and 97% in the lab ($p < 0.01$, $p = 0.93$, $p = 0.03$). Internet businesspeople chose the selfish option 61% of the time, which is not significantly different from students in the field ($p = 0.38$). In one treatment, ($6 for myself, $5 for the other) versus ($5 for myself, $15 for the other person), internet business people were significantly more altruistic than students in the field ($p = 0.05$).

In the second stage of the trust game (Game 1), students in the field chose ROLL (the trustworthy/altruistic option) 59% of the time compared with 33% of students in the lab, a difference which is highly significant ($z = 2.91$, $p < 0.01$). The level of altruism among students in the field at 59% was not significantly different from that among internet businesspeople whose level was 52% ($z = 0.75$, $p = 0.45$).

R6. Students in the field are no more trusting than students in the lab. Internet business people are more trusting than students.

In the first stage of Game 1, students in the field chose the trusting option 50% of the time compared with students in the lab who chose the trusting option 38% of the time ($p = 0.19$). Compared with internet professionals, who trusted 67% of the time, students in the lab are still significantly less trusting ($p = 0.05$). Unlike students in the lab, the lack of trusting behavior does not stem from pessimistic beliefs about whether trust will be repaid with kindness. Indeed, 51% of students in the field believe in the trustworthiness of the other person, this is comparable to internet business people ($z = 0.11$, $p = 0.91$) while significantly higher than the beliefs of students in the lab ($z = 2.32$, $p = 0.02$). Beliefs are well-calibrated—there is no statistically different difference between the beliefs about and actual trustworthiness of students in the field ($z = 0.94$, $p = 0.35$).

Table A1 in the appendix reproduces Table 2 above with the addition of the data from students in the field. As that table shows, results R5 and R6 are robust to the controls for gender and beliefs.

R7. Students in the field are less likely to lie than students in the lab. Internet business people lie about as often as students in the field.

As above, we use our preferred definition of lying of attempting to mislead the other person for financial gain. Under this definition students in the field chose to lie 52% compared to 71% of the time in the lab. This difference is highly significant ($z = 2.40$, $p = 0.02$). Students in the field lie more than internet businesspeople (46%), but the difference is not significant ($p = 0.47$). The results are similar if we use a more conservative definition of lying, counting only the subjects who believed that others would follow their advice. Under the conservative definition, students in the field lied 50% of the time compared to 70% of the time for students in the lab (and 45% of

the time for internet professionals). Again, there is no significant difference in the rates of lying by students in the field and internet business people while the difference between students in the field and students in the lab is now only marginally significant ($p = 0.10$) because there are fewer observations.

R8. Students in the field show similar effects to shame as students in the lab. Internet business are less driven by extrinsic motives (shame) than students in the lab or the field.

As before, we analyze the effect of our shame treatment (whether the choice of the second player was disclosed or not to the first player). This increases trustworthiness from 56% to 60% for students in the field. Students in the lab displayed a similar increase—going from 31% to 35%. The increase in trustworthiness for students in the field is not statistically significant ($p = 0.71$), as with students in the lab. The story is similar with regard to the effect of shame on trusting behavior. Students in the field increase trust from 49% to 51% compared to students in the lab who increase trust from 35% to 42% under the shame treatment. The increase in trust under the shame treatment is not significant for students in the field ($p = 0.83$).

Table 4 presents the results of a regression comparing students' response to shame with the response of internet business people. In the first column, we regress trustworthy behavior on an interaction of the shame treatment and whether the subject is an internet professional.[17] As Table 4 illustrates, there is a significant crowding out effect of shame on the trustworthiness for internet business people, but no such effect among students. The second column of Table 4 performs the same analysis using trusting behavior as the variable to be explained. Again, shame has a crowding out effect on internet business people; however, in this case, it does not rise to statistical significance.

Summary

To summarize, field settings appear to produce more pro-social behavior than laboratory settings. In two of the four measures of pro-sociality, populations of Berkeley students are more pro-social when playing these games in the field compared to playing the identical games in the lab. However, this alone does not account for the pro-sociality of internet business people. Under the hypothesis that internet business people are significantly *less* pro-social than students and that the results of our first set of experiments are purely the product of context, we would have expected to find students in the field being significantly more pro-social than internet business people. If anything, we find slight evidence in favor of the opposite conclusion: Internet business people are similarly altruistic and lying averse as students in the field. They are also more trusting and good behavior is less driven by extrinsic factors such as shame. Taken together, this supports the view that the pro-sociality observed in laboratory experiments is not selected out in populations of individuals in competitive industries.

[17] Since there are no significant differences across students in the lab and the field, we pool these populations.

7 Conclusion

The internet has widely been perceived as the "Wild West" in terms of business practices. Much has been written about unsavory behavior ranging from identity theft, Nigerian email scams, cybersquatting, and the ready availability of pornography to minors. The ease with which one can operate anonymously and vanish into cyberspace is thought to be a key factor allowing such practices to thrive. With this background in mind, we investigated the pro-sociality of domain traders and online adult industry professionals in a field setting using standard and novel laboratory experiments. We then ran these same experiments on students at Berkeley.

The external validity of pro-social behavior exhibited by students has often been called into question. The usual argument is that individuals with such preferences would be weeded out in highly competitive settings like the pornography industry. Yet, there was surprisingly sparse evidence to support this claim. While our priors were that internet business people would exhibit less pro-social behavior than students—indeed we expected that the purely self-interest homo economicus models would offer excellent predictions—we wanted to see the evidence for ourselves.

We were surprised by what we found. Across dimensions of trust, altruism, monitoring, and lying, internet business people were more pro-social than students. Moreover, the differences were not small. Compared to students in the lab, internet business people were twice as likely to both be trustworthy and over 50% more likely to trust in a trust game. Internet business people contributed over 250% more in dictator games. They lie one-third less often than students. The effects are more muted compared to students in the field, suggesting that field settings, in part, promote prosociality; however, even here, internet business people were, on the whole "nicer" than students.

The obvious follow-on question is why. Like the common wisdom, we think selection is responsible, but that it operates in precisely the opposite direction. Selfish business people find it harder to solve the "trust problem" and achieve success. This is especially true in the online space, where anonymity undermines the usual reputational strategies for achieving trust. In short, we think successful business people are more likely to be pro-social than not. The same pressures do not operate on students to the same degree. While we believe this is what is going on, our experiments were not designed to test this ex post explanation against any of several alternatives. Thus, our results should be viewed as more of a first word than a last. Investigating pro-social selection pressures in adult populations strikes us as an extremely promising, and important, path.

Appendix

Proof of Proposition 1. We consider the following strategy profile and show that it constitutes a PBE:

Pro-social individuals: Cooperate in the first period. In second period, if neither player in the match is revealed to have defected, choose Defect. Otherwise, choose Don't Trade.

Selfish individuals: Defect in first period. In second period, if neither player in the match is revealed to have defected, choose Defect. Otherwise, chose Don't Trade.

As usual, we solve using backwards induction. In the second period, if neither player has been revealed to have defected, since mutual defection comprises a Nash equilibrium, neither player can benefit by deviating. Likewise, for the case where either player has been revealed to have defected, since eschewing trade comprises a Nash equilibrium, neither player can benefit by deviating. Notice that both mutual defection and mutual non-trading are strict Nash equilibria as well.

Now consider the first period payoffs. Along the equilibrium path, a player earns

$$U(C) = (1-\pi)b + \pi d + (1-\pi r)c \tag{1}$$

from cooperating and

$$U(D) = (1-\pi)(a-\theta_i) + \pi c + (1-r)(1-\pi r)c$$

from defecting.

We will first show that neither type of player can benefit by deviating to no trade in the first period. The payoff under this deviation is $U(NT) = (1-\pi r)c$. Assumption A2 implies that, for pro-social types, $U(C) > U(NT)$, so this is not a profitable deviation. For selfish types, we require

$$(1-\pi)a + \pi c - r(1-\pi r)c > 0$$

which holds since $a > c$.

Next, we compare cooperating and defecting in the first period. It may be readily shown that cooperation is preferred to defection if and only if $\theta_i > \bar{\theta}$, where

$$\bar{\theta} = \frac{1}{1-\pi}\left(cr(\pi r - 1) + (a-b)(1-\pi) + \pi(c-d)\right)$$

By A3, pro-social individuals have $\theta_i = \theta > \bar{\theta}$, and thus pro-social individuals choose to cooperate. Selfish individuals will choose to defect if and only if $\bar{\theta} > 0$. Showing that $\bar{\theta} > 0$ is the same as showing the condition $cr(1-\pi r) + \pi(c-d) < (a-b)(1-\pi)$. We prove this condition as follows:

$$
\begin{aligned}
& (rc(1-\pi r) - \pi (c-d)) \\
< \; & (a-b)(1-\pi r) - \pi (c-d) \\
< \; & (a-b)(1-\pi-\pi r) \\
< \; & (a-b)(1-\pi)
\end{aligned}
$$

with both the first and second inequalities following from A1. Therefore, we have shown that pro-social individuals prefer to Cooperate whereas selfish individuals prefer to Defect in the first period.

Finally, we need to establish that the parameter conditions for θ are non-vacuous or, equivalently, that $\bar{\theta} < a - b$. This is equivalent to the condition that $cr(\pi r - 1) + \pi (c - d) < 0$. Notice that this expression is increasing in π and has a root when $\pi = \frac{cr}{c-d+cr^2}$. A3 then implies the desired inequality. This completes the proof. ■

Proof of Proposition 2. Owing to A1, there is no PBE where selfish types cooperate in the first period. Thus, the equilibrium above is the best equilibrium in which there is "exclusion" (no trade following some history). The welfare-maximizing non-exclusive equilibrium is where all types defect in every period. In that case, welfare is simply $2c$. Welfare in the equilibrium described in Proposition 1 is at least $U(C)$ given in equation 1. Notice that this expression is decreasing in π. Moreover $U(C) = 2c$ if and only if $\pi = \frac{b-c}{b-d+rc}$. However, by A2, π is smaller than this amount and hence welfare is higher under the equilibrium identified in Proposition 1. ■

References

Andreoni, James and Lise Vesterlund, "Which Is The Fair Sex? Gender Differences In Altruism," *The Quarterly Journal of Economics*, February 2001, *116* (1), 293–312.

Ariely, Dan, Anat Bracha, and Stephan Meier, "Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially," *American Economic Review*, March 2009, *99* (1), 544–55.

Benz, Matthias and Stephan Meier, "Do People Behave in Experiments as in the Field? Evidence from donations," *Experimental Economics*, September 2008, *11* (3), 268–281.

Camerer, Colin, *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton, NJ: Princeton University Press, 2003.

Cardenas, Juan Camilo and Jeffrey Carpenter, "Behavioural Development Economics: Lessons from Field Labs in the Developing World," *The Journal of Development Studies*, 2008, *44* (3), 311–338.

Eckel, Catherine C and Philip J Grossman, "Are Women Less Selfish Than Men? Evidence from Dictator Experiments," *Economic Journal*, May 1998, *108* (448), 726–35.

Edelman, Benjamin, "Large-Scale Registration of Domains with Typographical Errors."," 2003a. Harvard University.

__ , "Domain Name Typosquatter Still Generating Millions."," 2003b. Harvard University.

Edelman, Benjamin, "Red Light States: Who Buys Online Adult Entertainment?," *Journal of Economic Perspectives*, 2009, *23* (1), 209–220.

Fehr, Ernst and John A. List, "The Hidden Costs and Returns of Incentives-Trust and Trustworthiness Among CEOs," *Journal of the European Economic Association*, 09 2004, *2* (5), 743–771.

Gneezy, Uri, "Deception: The Role of Consequences," *American Economic Review*, March 2005, *95* (1), 384–394.

__ **and Aldo Rustichini**, "A Fine is a Price," *Journal of Legal Studies*, January 2000, *29* (1), 1–17.

Goldman, Adam, "Domain names: 21st century real estate," *Washington Post*, 2007. 7/22/2007.

Hopkins, Jim, "'Penthouse' makes $500M hookup with social site Various," *USA Today*, 2007. 12/12/2007.

Jenkins, Mark, *Beyond Tolerance: Child Pornography on the Internet*, New York, NY: New York University Press, 2001.

Joyce, Berg, Dickhaut John, and McCabe Kevin, "Trust, Reciprocity, and Social History," *Games and Economic Behavior*, July 1995, *10* (1), 122–142.

Karlan, Dean S., "Using Experimental Economics to Measure Social Capital and Predict Financial Decisions," *American Economic Review*, December 2005, *95* (5), 1688–1699.

Kendall, Todd, "Pornography, Rape, and the Internet," 2006. Working paper, Clemson University.

Kesmodel, David, *The Domain Game*, Xlibris Publishing, 2008.

Lazear, Edward, *Personnel Economics for Managers*, Wiley Publishing, 1998.

Levitt, Steven D. and John A. List, "What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World?," *Journal of Economic Perspectives*, Spring 2007, *21* (2), 153–174.

McCarthy, Kieren, *Sex.com: One Domain, Two Men, Twelve Years and the Brutal Battle for the Jewel in the Internet's Crown*, London, England: Quercus Publishing, 2007.

Ropelato, Jerry, "Internet Pornography Statistics," http://internet-filter-review.toptenreviews.com/internet-pornographystatistics.html 2006.

Roth, Alvin E., Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir, "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study," *American Economic Review*, December 1991, *81* (5), 1068–95.

Schechter, Laura, "Theft, Gift-Giving, and Trustworthiness: Honesty Is Its Own Reward in Rural Paraguay," *American Economic Review*, December 2007, *97* (5), 1560–1582.

Slonim, Robert and Alvin E. Roth, "Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic," *Econometrica*, May 1998, *66* (3), 569–596.

Tadelis, Steven, "The Power of Shame and the Rationality of Trust," 2009. Working paper, UC Berkeley.

Zillmann, Dolf and Jennings Bryant, "Pornography, Sexual Callousness, and the Trivialization of Rape," *Journal of Communication*, 1992, pp. 10–21.

Figure 1: Giving Percentage by Endowment In Game 3

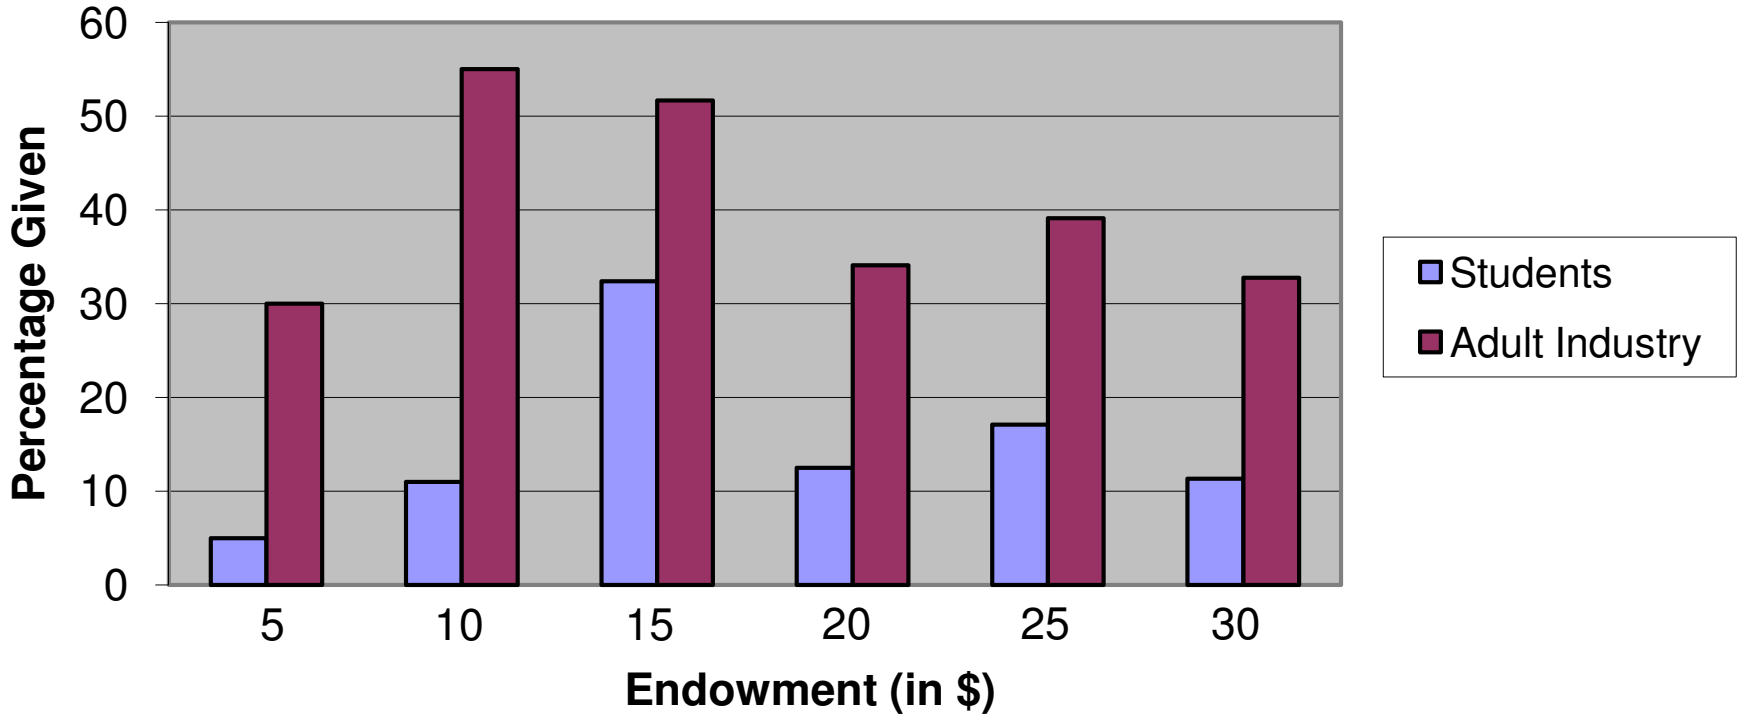

Figure 2: Lying Rates for Students and Internet Professionals

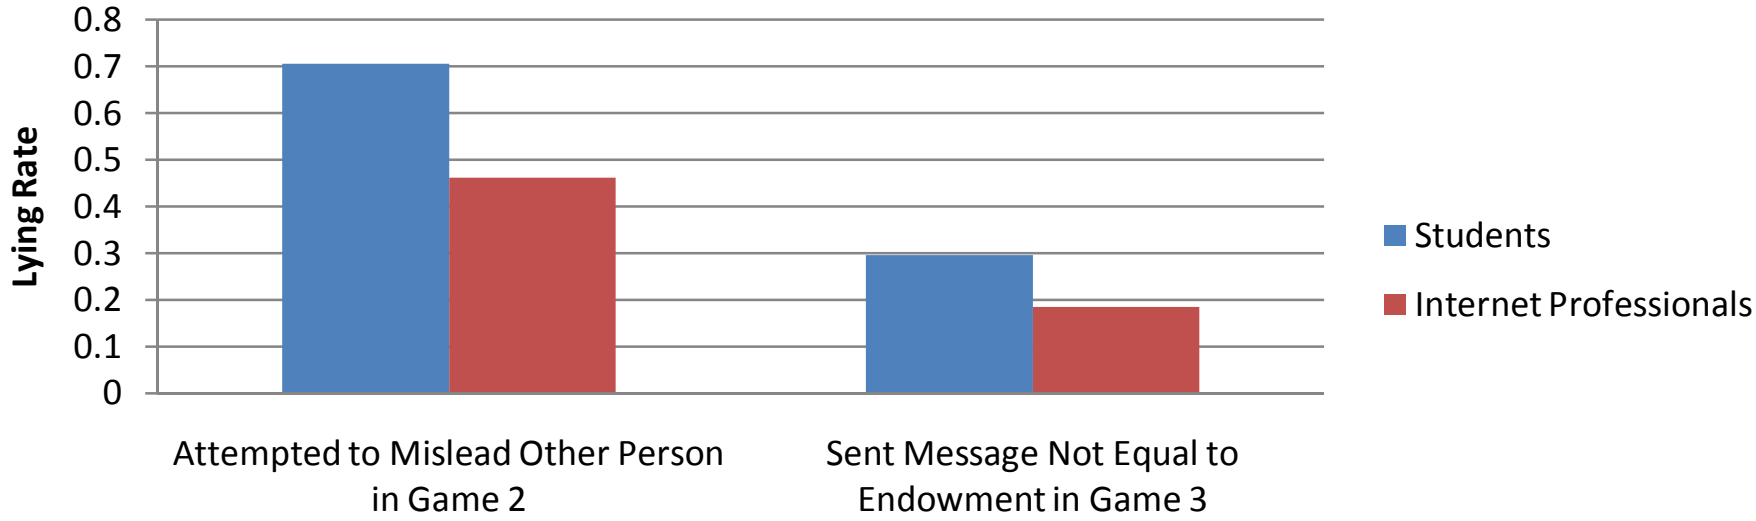

Table 1: Summary Statistics

	Domain Traders			Adult Professionals		
Variable	**Obs**	**Mean**	**Std. Dev.**	**Obs**	**Mean**	**Std. Dev.**
Domain investor / Entrepreneur	44	0.48	0.51	52	0.46	0.50
Industry Professional	44	0.41	0.50	52	0.27	0.45
Years involved with industry	43	5.60	4.43	51	5.50	4.58
High school or less	44	0.00	0.00	52	0.06	0.24
Some college	44	0.11	0.32	52	0.23	0.43
College	44	0.43	0.50	52	0.56	0.50
Graduate School	44	0.45	0.50	52	0.15	0.36
Total income less than 100k	39	0.26	0.44	49	0.41	0.50
Total income between 100k and 300k	39	0.41	0.50	49	0.49	0.51
Total income more than 300k	39	0.33	0.48	49	0.10	0.31
Asian	40	0.17	0.38	52	0.08	0.27
White	40	0.68	0.47	52	0.73	0.45
Male	43	0.86	0.35	52	0.69	0.47

Note: The year involved for adult professionals is years involved with the online adult entertainment industry. The mean number of years in adult entertainment as a whole is 5.95 years (s.d. 4.91)

Table 2: Share Given, Trustworthiness, and Trust

Dep. Var.:	**Share Given of One's Endowment (from Game #3)**	**Trustworthy (from Game #1)**	**Trust (from Game #1)**
	(1)	(2)	(3)
Internet Professional	0.252	0.313	0.236
	(0.053)***	(0.113)***	(0.109)**
Male	-0.056	-0.307	-0.124
	(0.054)	(0.114)***	(0.108)
Believe that Other Person Will be Trustworthy			0.404
			(0.096)***
Constant	0.173	0.439	0.306
	(0.046)***	(0.078)***	(0.080)***
Observations	106	95	95
R-squared	0.18	0.10	0.23

Notes: This table reports OLS regressions, with standard errors in parentheses. 'Trustworthy' equals 1 if the agent chose ROLL and 0 if the agent chose DON'T ROLL. `Trust' equals 1 if the agent chose IN and 0 if the agent chose OUT. (2) and (3) are very similar under a logit model.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 3: Game 1 Effects of Shame Treatment

	Trustworthiness		Trust	
	Shame	No Shame	Shame	No Shame
Domain Traders	35%	66%	61%	72%
Students	35%	31%	42%	35%

Table 4: The Impact of the Shame Treatment on Trustworthiness and Trust

Dep. Var.:	Trustworthiness	Trust
	(1)	(2)
Shame Treatment	0.046	0.047
	(0.086)	(0.085)
Internet Business Person	0.194	0.293
	(0.111)*	(0.110)***
Shame Treatment X Internet Business Person	-0.353	-0.163
	(0.164)**	(0.162)
Observations	186	186
R-squared	0.03	0.04

Notes: This table reports OLS regressions, with standard errors in parentheses. 'Trustworthy' equals 1 if the agent chose ROLL and 0 if the agent chose DON'T ROLL. `Trust' equals 1 if the agent chose IN and 0 if the agent chose OUT. In this table, internet business people are compared against both students in the lab and students in the field.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table A1: Trustworthiness and Trust

Dep. Var.:	**Trustworthy (from Game #1)**	**Trust (from Game #1)**
	(1)	(2)
Internet Professional	0.253	0.186
	(0.111)**	(0.108)*
Students in the Field	0.141	0.086
	(0.098)	(0.095)
Male	0.028	0.022
	(0.090)	(0.086)
Believe that Other Person Will be Trustworthy		0.308
		(0.080)***
Constant	0.374	0.282
	(0.076)***	(0.077)***
Observations	146	146
R-squared	0.05	0.14

Notes: This table reports OLS regressions, with standard errors in parentheses. 'Trustworthy' equals 1 if the agent chose ROLL and 0 if the agent chose DON'T ROLL. `Trust' equals 1 if the agent chose IN and 0 if the agent chose OUT. (1) and (2) are very similar under a logit model.

* significant at 10%; ** significant at 5%; *** significant at 1%