

**Trust among the Avatars: A Virtual World Experiment,
with and without Textual and Visual Cues**

by Stephen Atlas* and Louis Putterman**

Abstract. We invited “residents” of a virtual world who vary in real-world age and occupation to play a trust game with stakes comparable to “in world” wages. In different treatments, the lab wall was adorned with an emotively suggestive photograph, a suggestive text was added to the instructions, or both a photo and text were added. We find high levels of trust and reciprocity that appear still higher for non-student and older subjects. Variation of results by treatment suggests that both photographic and textual cues influenced the level of trust but not that of trustworthiness.

Keywords: trust, experiment, internet, virtual world, priming

JEL code numbers: C70, C91, D03

Forthcoming: Southern Economic Journal

* Marketing Ph.D. candidate, Columbia Business School, New York, NY 10027. saa2140@columbia.edu.

** Professor of Economics, Brown University, Providence, RI 02912. Louis_Putterman@brown.Edu.

Trust among the Avatars: A Virtual World Experiment, with and without Textual and Visual Cues

In recent decades, laboratory decision-making experiments have become a standard part of the economist's toolkit for testing theories of economic behavior, with the majority of experiments being conducted in university computer labs using students as subjects. While some results have been replicated using more diverse subject pools and while many replications suggest that inferences from student behaviors are not necessarily misleading, it is widely agreed that including more broadly representative subject pools in experiments would be desirable.¹

One reason most experiments continue to study students is that a properly incentivized experiment is costly, and most students already on a university campus have far lower opportunity costs of getting to an on-campus lab and participating in an experiment there than do individuals engaged in full-time employment in other locations. For experimenters, too, the costs of setting up experiments in other locations can be high relative to conducting them on campus. Thus, experimenting in university labs with student subjects carries substantially lower monetary and logistical costs.

The possibility of reaching a larger and broader subject population with modest logistical costs and cost per subject sparked our interest in conducting an exploratory experiment in a virtual world. Participants in the life-like on-line platform Second Life come from a broad demographic spectrum, and it has been estimated that about thirty thousand people are engaged in some activity in Second Life (including holding paid jobs, selling products, furnishing virtual homes, and socializing) at any given moment. We wanted to investigate the potential of this medium and population as venues for conducting economics experiments with diverse participants at reasonable cost.

The experiment we chose is the Trust Game introduced by Berg, Dickhaut and McCabe (1995, hereafter BDM) and subsequently replicated, with or without

¹ To be sure, experimental economics has come a long way since the days when nearly all experiments took place at universities. There are examples of experiments in which the subjects are fishermen, peasant farmers, pastoralists or even hunter-gatherers in developing countries (see, for example, Henrich *et al.*, 2001, Carter and Castillo, 2002, Barr, 2003, and Karlan, 2005) or truckers (Burks *et al.*, 2007) and businessmen (Fehr and List, 2004) in developed ones, and remarkable cross-country comparisons have been done (e.g., Herrmann, Thöni and Gächter, 2008). Nonetheless, a substantial majority of decision-making experiments in economics are still conducted at universities using students as subjects.

modification, in many settings (a partial survey is provided by Camerer, 2003). One reason for this choice is the game's simplicity: only two players need be matched, each makes a single decision, and the choices can be consecutive rather than simultaneous, so that no subject group as such need be assembled. Another reason is the game's intrinsic interest as a prominent illustration of how rational calculations intersect with social dispositions, and as an important tool for studying a problem lying at the heart of much economic interaction (Arrow, 1974). How the social element in the trust game will play out in the anonymity of the virtual world is a matter of both theoretical and practical interest, the latter owing to the fact that trust in e-commerce and virtual environments has commercial and other applications (Benbasat, Gefen and Pavlou, 2008; Chesney, Chuah and Hoffman, 2009).

In addition to conducting a standard trust game in order to compare results of an otherwise similar protocol with different environments and subject pools, we decided to experiment with our virtual lab by varying two dimensions and performing exercises in comparative statics. First, we varied the physical environment by placing one of two photographs, or none, on the virtual lab room's wall, following a recent psychology experiment suggesting strong susceptibility of social behaviors to subtle but economically irrelevant cues (Bateson, Nettle and Roberts (2006)). While this kind of environmental manipulation could also be done in a brick-and-mortar lab, ease of handling digital images in the virtual environment makes it particularly easy to implement in a manner that is salient while its treatment status remains unrecognized by subjects. Second, we varied the ideational environment by adding differently suggestive wordings at the conclusion of the instructions of two treatments.² While this is easy enough to do in a physical lab, we wanted to contrast the impact of words with that of pictures in the same setting. In both dimensions of variation, we contrasted a trust-encouraging to a trust-discouraging treatment and each to the baseline experiment which closely follows BDM. We also implemented treatments that include both visual and textual cues.

² An example of an experiment in which the impact of suggestive text is demonstrated is Dal Bó and Dal Bó (2009), who find that including certain passages as a form of "moral suasion" in the instructions of a public good game has transitory effects in the absence of punishment opportunities and lasting effects when there is an opportunity to punish.

Our virtual lab visually resembled a conventional one, with an entry area (see Figure 1) containing a receptionist, a portal (“teleporter”) to individual experiment rooms, and six such rooms each designed to comfortably seat one subject. The subject pool differs from that in most university experiments in having a much broader range of ages and occupations (we discuss problems of verification below). Subjects participated from any location allowing internet access, and did so from at least 39 different countries. Data collection was 100% automated and ran twenty-four hours a day with a round-the-world subject pool.

Potential problems are posed for experimenters by the fact that subjects visit the virtual lab in the guise of “avatars,” with information they provide about their real world demographics being difficult to verify. Possibly subjects might have their avatars act in ways that they themselves would not. Also, the same individual might play multiple times using different avatars, although this danger can be minimized (see below). The heightened anonymity of interactions among such subjects may have its advantages, however, in that a potentially worrisome desire to please the experimenter (so-called “experimenter demand effect”) is attenuated because no embodiment of the experimental team appears in the lab (the receptionist was an automaton, and clearly so). There is a potential drawback from the lack of hard evidence of counterpart authenticity,³ yet if subjects act, under such extraordinarily anonymous conditions, according to social norms and in the expectation that others also will do so, this would be strong evidence that such behaviors and beliefs are deep-seated.

Average choices by our first-movers are remarkably similar to those in BDM. In contrast, our second-movers are significantly more trustworthy, rendering trusting profitable (whereas it is slightly unprofitable on average, in BDM). Each of our photo and word manipulations shifts trust in the direction expected, illustrating the potential of the medium, while also introducing some new observations about the powers of different kinds of interventions (for instance, whereas the trust-discouraging cue has the larger

³ For internet-mediated (but not virtual world) experiments, Eckel and Wilson (2006) find that subjects’ doubts that there is a real counterpart affect first but not second-mover behaviors in a trust game. Lack of certainty that there is a real counterpart might help to explain lower first-mover sending in the internet-based treatment of Chesney *et al.* (2009) as compared with their virtual world treatment in which subjects could visually verify the presence of one another’s avatars in the virtual lab room. Conversely, evidence that the counterpart exists may be one factor explaining the greater trusting and trustworthiness in trust games with pre-play communication (Ben-Ner, Putterman and Ren, 2007).

impact among the verbal treatments, the trust-encouraging cue is the more effective of the visual ones). The amounts returned by second-movers are significantly correlated with the amounts they are sent but unaffected by treatment cues.

Since broadening the subject pool beyond university students is one of our interests, we also investigate the possible dependence of behaviors on demographic variables. We find that while age, student status, and income affect behaviors, those effects are not consistent with the suggestion that older, employed, or higher-earning individuals' behaviors should conform with the predictions of standard economic theory better than the behaviors of students do. On the contrary, our older, higher-earning, non-student subjects tend to be significantly more trusting and trustworthy than their younger, lower-earning and student counterparts.⁴ We also explore the possibility that individuals having different demographic characteristics might respond differently to given treatment variables, finding this to be for the most part untrue, but with interesting exceptions for gender and age range splits. And we confirm Glaeser, Laibson, Scheinkman and Soutter's (2000) finding that self-reported belief that people can be trusted predicts own trustworthiness but not trusting.

While ours is among the first generation of economics experiments in virtual worlds, a few others are also beginning to appear. Chesney, Chuah and Hoffmann (2007) report play of a number of familiar games including the dictator game, the ultimatum game, and the public goods game, and find subject behaviors closely resembling those in conventional labs. The subjects in Fiedler and Haruvy (forthcoming), Fiedler (2009) and Fiedler, Haruvy and Li (2008) play trust games following in-world communication with prospective partners, and the latter paper includes a control treatment without communication.⁵ The subjects in the VE condition in Chesney, Chuah and Hoffman (2009) play a version of the trust game in Second Life in sessions in which the avatars of potential counterparts sit in a large virtual lab room separated by modest partitions. In contrast with the latter papers, our experiment follows BDM in having complete anonymity and isolation of first and second movers. And ours is (to our knowledge) the

⁴ To be sure, the additional trusting can be justified *ex post* as being fully rational and self-interested, given that second-mover behaviors render trusting profitable in our data, but the same cannot be said of second-movers' greater trustworthiness.

⁵ Castronova *et al.* (2008) study how demand for a virtual product varies with its price but do not engage subjects in a laboratory-style decision-making experiment.

first instance of a persistent economic experiment in a virtual world that automates recruitment, instruction, measurement and compensation, eliminating any direct interaction with the experimenter. Finally, our experiment uses payoffs scaled closer to earnings from in-world jobs than to payoffs in most laboratory experiments. Despite this, the large majority of our subjects report that the prospect of earning money was an important motivation for participating, and almost all agree that the range of earnings was reasonable.

The Experiment

Virtual worlds are persistent, online communities in which individuals interact with the synthetic environment and each other through avatars embodying their virtual selves. These highly customizable spaces are tailored to service shared objectives requiring a high degree of interactivity between geographically dispersed individuals. Some worlds are tailored to servicing institutional objectives. Educators use virtual worlds for distance learning and immersive simulation. Business users have utilized virtual worlds for personnel recruitment, trade conferences, to collaborate across distances and to expand brand awareness. Virtual worlds have enabled chronically or temporarily disabled individuals to overcome obstacles to socialization. Recreational virtual world users play socially interactive games and co-create shared realities, finding simulated experiences to be more engaging and satisfying than other sedentary hobbies.

For our experiment, we chose the world Second Life due to its combination of a potential subject pool conservatively estimated in the hundreds of thousands,⁶ control afforded with respect to the configuration and appearance of the local virtual environment, relatively realistic and non-violent ambiance, and the possibility of legally purchasing virtual currency with which to pay subjects “in world.” The experiment was conducted in a second generation version of a virtual experimental laboratory that was constructed under our direction on Second Life’s Etopia Island, incorporating improvements over the version used in Atlas (2008a).

⁶ Approximately 1.45 million “residents” logged in to Second Life between October 25 and December 24, 2008. A resident is defined by Linden Lab as a unique avatar, and it is noted that users may possess more than one avatar, rendering the unique user count lower. At the time of the experiment, an average of 48,000 users were in-world at any given moment.

Most subjects learned of the experiment opportunity by reading a classified ad in Second Life and choosing to “teleport” to the lab entrance.⁷ Upon entering, they found their avatar in a lobby where it encountered an automated, human-appearing greeter seated at a desk, who (when the subject clicked as instructed by a sign) invited their avatar to step onto another “teleporter” from which they emerged into one of six lab rooms each of which contained a single chair, an initial instruction sign, and in some treatments a framed photograph hanging on the wall. Detailed instructions appeared in a “blue box” on the upper right of the subject’s screen. Only one subject at a time was sent to a given lab room, and no local communication with other players was permitted.⁸ The same avatar could not participate twice, and a variable lag in timing between first-mover and second-mover participation (for any given pair) assured that subjects could not arrange to play with either themselves (under an alternate avatar) or a friend.⁹

The experiment followed closely the structure of the trust or investment game introduced in BDM, where two subjects who remain anonymous to one another are each given an endowment of ten one dollar bills. The first-mover in that design, whom we’ll call A, has to choose an integer amount $X_A \in (0, 1, \dots, 10)$ to send to the second-mover, B. As known in advance by both subjects, exactly three times the amount of money sent by A is delivered by the experimenter to B. B may then send to A any or no part of this tripled amount, i.e. $0 \leq X_B \leq 3X_A$ (again, confined to integer values), which is not subject to further multiplication. Thus, an amount that can vary between \$20 and \$40 (the total being determined solely by A’s choice) is to be divided between A and B in a manner determined partly (if $0 < X_A < 10$) or entirely (if $X_A = 10$) by B. Equilibrium play assuming common knowledge of rationality and self-interested preferences is that A

⁷ According to their responses to questions asked after the decision task, 67% of subjects learned about the experiment by reading the ad, 15% by word of mouth, and 18% by virtue of their avatar happening to pass by the lab building, over the door of which was a sign reading “Experimental Economics Lab; Participate in a 15-Minute Study; Earn L\$100 – 400.”

⁸ We checked for local communication subsequently and were prepared to drop observations in which the requirement was violated, but found no relevant instances.

⁹ Two individuals could try to be matched with one another, but there were very low odds that they would succeed, and (equally important) no way they could obtain confirmation that they had done so, even if they were separately in contact by phone, e-mail, or other means. See relevant wording of instructions, in the Appendix. More details on construction of the experiment lab, choice of Etopia island, and criticisms received on the initial design are found in Atlas (2008b) and discussed in Bloomfield and Rennekamp (2009), while improvements made are discussed in the Appendix, which also includes wording of the advertisement and script of the greeter’s message.

sends nothing since B will return nothing. In BDM and subsequent replications, however, the large majority of first-movers are found to send positive amounts, and a majority of second-movers receiving money also return some, with average outcomes having B's earning unambiguously more and A's either more or less than in the predicted equilibrium, depending upon the trial.

In our neutral or BASELINE treatment, we replicated BDM's procedure except for those logistical details and instructions specific to the handling of physical currency and the protection of double-blind anonymity through passing of marked envelopes, assignment of room monitors, use of mailboxes, etc.¹⁰ With these exceptions, procedures were explained to subjects in as nearly as possible the same terms as used by BDM, and each knew the full instructions given to subjects in both roles before being informed of the role that he or she would be assigned to.¹¹ We substituted endowments of L\$100, where L\$ stands for Second Life "Linden" dollar, for the BDM endowments of (U.S.) \$10, and we allowed subjects to send and return L\$ integer amounts.¹² While at the time the experiments were conducted, L\$100 could be purchased for considerably less than \$10, our average subject earned more in twenty minutes of our experiment than in an hour of their Second Life employment, and 93% of subjects strongly agreed or agreed that the range of possible earnings from the experiment was reasonable.¹³

¹⁰ Note that in our design, saying that 'A sends money to B' or 'B sends money back to A' is technically accurate only in the sense of money credited towards later payment, which takes place at the end of the experiment (an electronic payment in Second Life taking place within minutes of B's decision). One cannot rule out that both amounts sent and amounts sent back in BDM were influenced by the fact that the subjects in question were parting with physical dollar bills that they were free to instead place in their own pockets. However, many other experiments using the basic BDM design have been conducted by computer or with passing only of decision sheets, with money payment at the end of the session, as is the rule in economics experiments. No clear difference in behaviors is apparent in the trust game experiments conducted with physical bills and those with payment at the end of the interaction.

¹¹ Unlike our experiment and BDM, Chesney *et al.* (2009) paid their subjects show-up fees separate from the game payoffs and did not match the initial L\$1000 given to A with any such grant to B. Unlike our experiment and BDM, in Fiedler and Haruvy (forthcoming) and Fiedler *et al.* (2008) B players could send their counterparts not only any portion of their tripled receipts but also their initial endowment (which as in BDM and our experiment, is the same as the endowment of A players).

¹² Having choices be over the integers (as in BDM) rather than in increments of ten seemed more natural to us. That it causes the action space to be more finely partitioned, with 101 rather than the 11 choices available to first-movers as in BDM, is a by-product of this decision that we don't consider to be of particular importance.

¹³ At the time of our experiment, 100 Linden dollars could be purchased from the LindeX™, the official Linden dollar exchange, at a cost of \$0.38. The typical subject devoted less than twenty minutes to our experiment, compared to 60 to 90 minutes for subjects in BDM. 55% of our subjects report holding a job in Second Life, with hourly wage rates averaging L\$172.77 an hour (USD\$0.65 per hour), approximately

We randomly assigned subjects, who had no knowledge that multiple treatments existed,¹⁴ to either our BASELINE treatment or to one of six others. The other treatments differed from BASELINE either by the presence of one of two photographs on the virtual laboratory wall, by the addition of a pair of sentences to the end of the instructions, or both.

Evolutionary psychologists argue that human sociality is governed by innate predispositions that are responsive to environmental triggers.¹⁵ A dramatic recent example is reported by Bateson, Nettle and Roberts (2006), who show that honor system payments for coffee, tea and milk were almost three times higher when a narrow strip showing a photograph of human eyes directed at the user was included on the price list above the payment box in a psychology department's coffee room as when the same strip contained photographs of flowers, even though no user reported noticing the pictures. Because a large proportion of subject behaviors in trust games depart from the prediction based on strict rationality and self-interest, it seemed likely to us that decisions to trust or be trustworthy might be sensitive to similar cues. Also, ease of manipulation of visual fields may be a particular advantage of virtual world experiments.¹⁶ We looked for contrasting photographs that might, on the one hand, promote an optimistic social

one twentieth of the same individuals' average self-reported wage in real world employment. Residents use the currency to purchase items such as virtual clothing and accessories, to pay rent on stores and virtual apartments, and to buy services from other players. Examples of in-world prices include outfits sold for L\$100, hair styles for L\$100 to L\$200, and weekly rents on shop spaces and houses in a similar range. While the range of earnings possible in the roughly contemporaneous experiments of Chesney *et al.* and Fiedler and co-authors were far higher, we believe that with about a million "residents" in Second Life, few if any of our subjects knew of their roughly contemporaneous experiments, which we ourselves were not aware of.

¹⁴ A possible exception would be subjects who happened to encounter previous subjects either on their way out of the lab or subsequently, and who furthermore discussed relevant details of the experiment with them. Since only 15% of subjects reported learning of the experiment by word-of-mouth, and since very detailed discussions and recollection seem unlikely, we think it safe to assume that few if any subjects were aware of treatment differences.

¹⁵ An overview is provided in Buss (2008), a clear popular treatment in Wright (1994). For a discussion from the standpoint of economics, see Hoffman, McCabe and Smith (1998). Social scientists taking a more "co-evolutionary" view, for example Boyd and Richerson (1985), give more equal emphasis to social and cultural evolution, with the environment shaping the behavioral expression of genetic predispositions over time, but such views also allow for innate social predispositions (see Ben-Ner and Putterman, 1998).

¹⁶ One could vary what pictures hang on the wall of a brick-and-mortar lab room also. However, the virtual environment makes it easier to experiment with different placements, sizes, and lighting, to rule out the possibility of prospective subjects noticing changes being made, and to assign subjects to rooms having one or another photo or none with no possibility of seeing the other rooms (thanks to travel by "teleportation"). More importantly, perhaps, the specific visual manipulations we tried only scratch the surface of what is possible in this medium, so that our treatments might be compared to taking a new vehicle for a quite modest test drive.

outlook, and on the other, a more cautious and self-interested (or self-protective) approach. The upper portion of Figure 2 shows the photos selected for what we'll call the OPTIMISTIC and CAUTIONARY treatments, respectively, while the lower portions show experiment rooms as viewed by a subject under the respective treatment (in non-photo treatments including BASELINE, the room was identical but for the absence of both picture and frame).

Although manipulation of instruction text can as easily be done in a physical as in a virtual lab, we were interested in comparing the impact of words to that of pictures holding all other design elements constant. We chose textual passages with the parallel aim of encouraging either optimistically social or cautious and self-interested play, calling the treatments in question COOPERATIVE and COMPETITIVE. The wordings are:

[COOPERATIVE treatment] By working together and acting fairly, both participants can double their money. (The first decision-maker can send L\$100, the second decision-maker can return L\$200, and they therefore each earn L\$200 instead of each earning 100).

[COMPETITIVE treatment] However, assuming that second person wants to earn as much as possible, he/she will send none of the money received back to the first person. So if the first person sends L\$100 and the second person keeps the largest amount of money, he/she will earn L\$400 and the first person will earn zero.

The full instructions are shown in the Appendix.

We expected at the outset that our visual and textual manipulations would work through different pathways, the former being automatic and unconscious, the latter infusing the contents of deliberative thought.¹⁷ Their effects might have differences of

¹⁷ The distinction between conscious and automatic mental processes is discussed by Bargh and Chartrand (1999). Burnham, McCabe and Smith's (2000) example of trust games using alternatively the terms "partner" versus "opponent" and Ross and Ward's (1996) example of subjects recruited to play prisoners' dilemma games described alternatively as "the Wall Street Game" and "the Community Game" provide

quality or degree, and would not necessarily be additive.¹⁸ Nevertheless, to check whether the visual and textual cues had reinforcing effects, we also implemented a COOP-OPTIMISTIC treatment in which subjects were exposed simultaneously to the OPTIMISTIC visual cue and the COOPERATIVE textual cue, and a counterpart COMP-CAUTIONARY treatment.

We attempted to collect twenty-five valid observations of paired first- and second-movers interacting in each of the seven treatments, although numbers exceeded this target in one case due to an initial difficulty in the randomizing of treatment assignments.¹⁹ We then analyzed behaviors using both the full samples of observations and samples from which we eliminated observations potentially attributable to the same individual playing under two identities, as indicated by a common IP address.²⁰

Results: first-mover sending

illustrations of how changing the description of a decision task can alter subject behaviors. Our added text treatments bear some resemblance to these manipulations but represent attempts to introduce more conscious lines of thought in subjects' minds. In addition to the example by Bateson, Nettle and Roberts (2006), described in the text, other examples of priming by photograph and film, respectively, are found in Scharlemann, Eckel, Kacelnik and Wilson (2001) and Kirschsteiger, Rigotti and Rustichini (2006). On the psychology of priming, see Klauer (1997). Page and Putterman (2010) contrast a trust game described to subjects using the standard terms including "send" and "send back," as initiated by BDM, and a formally equivalent game presented in entirely spatial terms without mention of such terms.

¹⁸ Note also that we would not want to generalize from findings from our specific textual treatments, with their explicit discussion of the strategy of play, to a larger universe of text-adding treatments, which could use suggestive terminology, employ words to achieve more subliminal priming effects, or embody still other approaches.

¹⁹ Subjects were initially assigned to a treatment through an in-world randomization script. However, the randomization caused some treatments to be favored slightly. Accordingly, we switched to having treatments be assigned by rotation, to more easily attain roughly equal numbers. Small differences in completion rates and in numbers of duplicate IP subjects account for remaining differences in numbers of valid observations.

²⁰ We attempted to minimize the possibility of individuals participating multiple times first by simply telling subjects that this was not acceptable as "the same person repeating the experiment compromises the integrity of the study," secondly by making sure a given avatar does not play twice, and third by instituting a procedure to capture IP addresses. We think few people would go to the trouble of changing avatars and IP addresses in order to participate twice in a study with rather modest returns. Two observations sharing the same IP address could indicate that the same person is playing twice, perhaps using different avatars, or that roommates or members of the same family played the game, conceivably sharing information (for instance, comparing one another's treatment cues). In our restricted samples, we eliminate all of the observations for which the IP address is not unique, not merely the later-occurring observations from a given address. 14.6% of observations are associated with a non-unique address, with the proportion of such dropped observations ranging from 2% (in the CAUTIONARY treatment) to 26% (in the COMPETITIVE treatment). On the trade-off between more detailed verification procedures and privacy and brevity, see also our discussion in footnote 30.

Table 1a shows the average amount sent by subjects in the A role by treatment. It reports the non-duplicate-IP observations, the full sample averages being similar. Considering first behaviors in the BASELINE treatment and average behaviors in the experiment as a whole, we note that the proportion of endowment sent by BASELINE treatment subjects lies only slightly below the average proportion sent in all treatments combined, 51.8%, and that this is remarkably similar to the proportion sent by BDM's subjects, 51.6%. It likewise resembles the amounts sent in other one-shot trust game experiments having the same payoff structure, such as Carter and Castillo (2002) (53.0%), Schechter (2007) (46.8%), and Danielson and Holm (2005) (52.2%), as well as the proportion of a L\$1000 endowment sent by Fiedler *et al.*'s control treatment subjects (53.4%).²¹ There is no indication here that smaller stakes, a more diverse subject pool, or a virtual environment alter average first-mover behavior in our experiment.

Turning to differences among treatments, Table 1a shows averages ranging from a low of 33.1% of the endowment in the COMPETITIVE textual cue treatment to a high of 70.4% of endowment in the OPTIMISTIC visual cue treatment. Average sending in the BASELINE treatment, with neither visual nor textual cues added, lies near the middle of this range, at 49.3%, with averages of the other treatments each falling above or below the BASELINE average in the directions expected. The difference in sending between the OPTIMISTIC and CAUTIONARY visual treatments and those between the COOPERATIVE and COMPETITIVE textual treatments are large: a spread of L\$21.54 (21.5% of endowment) between the two visual treatments and a difference of L\$26.01 (26.0% of endowment) between the two textual treatments, both differences being

²¹ We focus on the most similar comparison treatments, ignoring variants in which the second-mover is given no endowment or is permitted to return both the endowment and money received from the first-mover, variants in which the sent amount is doubled rather than tripled, ones with multiple interactions, binary versions of the game, etc. We consider Haruvy, Fiedler and Li's Control 3 treatment in which there was no pre-play communication, avatars were seated in different virtual rooms, and amounts sent by first-movers were tripled. Chesney *et al.* (2009)'s virtual world treatment first-movers sent a relatively high 61.2% of their endowments. This more trusting behavior might be attributable to the fact that six to eight subjects' avatars played simultaneously in the same room, so that the avatars of one's potential counterparts were "socially present," an intentional design element driven by their goal of comparing virtual world interactions to the less personal interactions characteristic of electronic commerce. Recalling that their second-movers did not receive the L\$1000 endowment that the experimenters gave to their first-movers (although both received L\$500 show-up fees), it is worth noting that some sending may have been motivated by fairness rather than trust (see Cox, 2004).

significant at the 5% level in two-tailed Mann-Whitney tests.²² The combinations of textual and visual cues fail to be mutually reinforcing in their effects, with the difference between COOP-OPTIMISTIC and COMP-CAUTIONARY sending averages being only L\$8.00 (8.0% of endowment) and not statistically significant. Each average nevertheless lies on the side of the BASELINE treatment's average on which the combined treatment would be expected to place it provided that combining treatments doesn't completely negate their individual effects.

When a Tobit regression is estimated for all observations with first-mover sending as dependent variable, treatment dummies as explanatory variables, and BASELINE as omitted category, only the OPTIMISTIC treatment dummy is statistically significant at the 5% level (see column 1 of Table 2). However, F-tests indicate that a number of the coefficients differ significantly from each other. In particular, differences between the coefficients for the main paired comparator treatments are significant: those of the COOPERATIVE and COMPETITIVE (textual) treatments differ at the 1% level, those of the OPTIMISTIC and CAUTIONARY (visual) treatments at the 5% level. The COMPETITIVE and OPTIMISTIC coefficients differ from each other at the 1% level of significance, as do the OPTIMISTIC and COMP-CAUTIONARY coefficients. The COMPETITIVE and COOP-OPTIMISTIC coefficients differ at the 5% level, and the COMPETITIVE and CAUTIONARY coefficients differ at the 10% level.²³

The effects of the two visual and the two textual treatment pairs differ in their “centering.” Relative to BASELINE, the OPTIMISTIC visual cue induces a dramatic upward shift in average sending whereas the CAUTIONARY visual cue induces only a statistically insignificant and economically small decline. Relative to BASELINE, the COOPERATIVE textual cue induces a smaller (although in this case more economically meaningful) increase in sending than the decrease induced by the COMPETITIVE cue.

²² Appendix Table A.1 provides a summary of Mann-Whitney test results for differences among all treatments both in amounts sent by first-movers and proportions returned by second-movers.

²³ Another way of observing these differences is to use a treatment pair member as the base. For example, when only the observations of the COOPERATIVE and COMPETITIVE treatments are included in a regression to predict first-mover sending, a dummy variable for COOPERATIVE treatment obtains a large positive coefficient (42.08) with $p = .014$. A parallel exercise using only the OPTIMISTIC and CAUTIONARY treatment observations yields a large positive coefficient (38.03) on the OPTIMISTIC treatment dummy with $p = .031$.

Looking at the comparatively weaker effects of the combined text-and-visual treatments, a larger sending difference relative to BASELINE is induced by COOP-OPTIMISTIC than by COMP-CAUTIONARY.²⁴ A possible explanation might be that a textual intervention has more power to move sending down than up because spelling out the logic of Nash equilibrium play between perfectly self-interested actors constitutes a bigger departure from many subjects' initial outlooks than does spelling out the appeal of cooperation. (That only a small minority send nothing in most trust game experiments suggests that the tendency to trust or to anticipate reciprocity is "natural" to most subjects, whereas the logic of sub-game perfect equilibrium play is not.) With respect to the visual cues, the photo of two cuddly pandas may have a more unambiguous power to imbue subjects with good feelings than the fish photo has to do the opposite because the social inclinations that lead to trusting in the BASELINE treatment are more readily reinforced through this emotive channel, whereas the caution subjects bring to their decisions is more susceptible to rational argument than to emotional reinforcement. Whatever the explanation for the asymmetry, it is worth noting that the COMPETITIVE and OPTIMISTIC treatments yielded differences in average first-mover sending of a magnitude such that sending in the latter averaged more than twice that in the former.

Decisions by second-movers

Column 5 of Table 1b shows, by treatment, the average proportion returned by subjects in the B role who received a nonzero amount from their counterpart in the A role (the only B's having a decision to make). The overall average works out to 47.3% of the tripled amount, with BASELINE treatment subjects returning an average of 50.7% of the amount received from their counterparts. These proportions are considerably higher than the 33.3% that causes the first-mover to break even, higher than the 29.8% average proportion returned by BDM's subjects, and above most other results in the literature, including BDM's "social history" treatment (37.1%) and Danielson and Holm's (2005) average over Tanzanian and Swedish subject pools (35.8%). The difference with our

²⁴ Non-parametric tests for differences with BASELINE sending yield results similar to those of the Tobit regression (Table 2, col. 1). In two-tailed Mann-Whitney tests, only sending in the OPTIMISTIC treatment is statistically different from that in BASELINE, in this case at the 10% level. The sending difference between COMPETITIVE and BASELINE falls a bit short of this level and achieves significance only in a one-tailed test. No other differences from BASELINE are statistically significant.

subjects' 47.4% average return proportion is smaller, however, in some studies: Burks, Carpenter and Verhoogen's (2003) base treatment (40.0%), the result in Schechter (2007) (43.7%), and the proportion returned by second-mover recipients of positive amounts in Fiedler *et al.*'s Control 3 treatment (41.7%), all showing that returns above the break-even level are not so uncommon.²⁵ The greater degree of reciprocity evinced by our subjects might be partly attributable to lower real-money stakes, but it is still remarkable in view of the importance subjects report that they attached to earning Lindens as a motivator of their participation, and especially given the unusually high degree of anonymity both among players and between players and experimenters. Moreover, it runs dramatically counter to the seemingly sensible intuition of at least one respected experimental economist that "[i]n the virtual world ... you're virtually free of ... reputational concerns, so you might get people acting in a more self-interested way."²⁶ Along with Chesney *et al.*'s average return rate of 38.5% of received amount in their virtual world treatment, but still more so in view of the greater anonymity of our design, our subjects' behaviors suggest that human sociality runs quite deep, being brought along by most subjects even as they manipulate cartoonish self-representations in a make-believe world.²⁷

Looking at the individual treatments and again restricting attention to cases in which second-movers had received positive amounts, OPTIMISTIC, COMP-CAUTIONARY, and COOP-OPTIMISTIC second-movers returned on average between 44 and 49% of the tripled amounts received by them, while CAUTIONARY and COMPETITIVE treatment subjects returned 42 and 39%, respectively, and COOPERATIVE treatment subjects were the only ones returning more than BASELINE counterparts, an average of 57%. Only two treatment differences in percentage returned are statistically significant in two-tailed Mann-Whitney tests: subjects in the COOPERATIVE verbal treatment returned more than those in both the COMPETITIVE verbal treatment and the CAUTIONARY visual treatment, significant at the 10% level.

²⁵ Second-movers return an average of 42.5% of the tripled amount in the no communication treatment in Ben-Ner, Putterman and Ren (2007), a result not included in the main comparison because subjects engaged in several interactions (each with a different partner).

²⁶ John List, quoted in Foster (2007).

²⁷ A similar interpretation might also be given to the considerable indications of reciprocity or fairness on the parts of participants in the uninterpreted spatial treatment of Page and Putterman (2010).

The ordering of compared percentages returned is consistent with expectations of which verbal and visual treatments would tend to elicit more cooperative and which more competitive behaviors both in the cases just mentioned and in those in which the differences lack statistical significance, with the exception that the proportion returned is slightly higher in the COMP-CAUTIONARY than in the COOP-OPTIMISTIC treatment.

These raw comparisons of proportion returned by second-movers are of more limited value than the comparisons of first-mover sending that preceded them because the proportion of tripled funds that second-movers return could be influenced by the amount received. Some studies, for instance Ben-Ner and Putterman (2009), find that the amount sent by the first-mover is a significant positive predictor of the proportion returned by the second-mover (“trustworthiness” rises with “trusting”), while others, including Cox (2004) and Brühlhart and Usunier (2004) find negative relationships. For our data, Tobit regressions in which the proportion returned by B is the dependent variable and the amount sent by A and treatment dummy variables are the explanatory variables, fail to show a significant effect of A’s sending on the proportion returned (see Table 3, column 1).²⁸ Nevertheless, controlling for amount sent alters conclusions about treatment effects on proportion return. Specifically, with this control included and with BASELINE treatment as default, none of the treatment dummies obtains a significant coefficient (see Table 3). Including interaction terms between amount received and treatments doesn’t change this outcome. Nor do any of the estimated treatment dummies differ significantly from each other according to F-tests. We conclude that while treatment differences significantly affect the choices of first-movers, they do not significantly affect those of second-movers once the behaviors of their first-mover counterparts are controlled for.

Note that failure to find a significant association between first-mover sending and second-mover return *proportion* does not mean that there is no correlation between sending and returning. A Tobit regression of *amount* (rather than proportion) returned by second-movers on amount sent by first-movers, with 142 observations, obtains a coefficient of 1.60 on the explanatory variable, with *t*-statistic 9.96. The numerical

²⁸ Only observations for those B’s who received positive amounts are included in our regression. While Fiedler et al. (2008) report a regression in which proportion returned is significantly positively correlated with amount sent, this is likely to be influenced by the inclusion of observations in which the first-mover sent zero.

average return on Linden dollars sent by first-movers, according to Table 1b, is 142% ($\approx .473/.333$), which falls short of the proportion needed to equalize the final earnings of B and A but nevertheless represents a full 42% return on trusting for first-movers. Thus a first-mover with perfect foresight regarding the distribution of second-mover responses, if intent on maximizing her earnings, should have sent her entire L\$100 endowment, regardless of treatment.²⁹

Subject characteristics, source of local currency, and their effects

A major goal of our experiment was to include older, non-student, and employed individuals in our sample at modest cost in terms of subject payments and logistics. But before we discuss our subjects' demographic characteristics to assess our success in this respect, we need to note that apart from IP addresses, our knowledge of subject demographics depends on information subjects provided in response to our post-decision demographic questions. We have no way of ruling out that any given response might be fictitious, so those of our findings that depend on this portion of our data must be treated with caution. Since we provided strong assurances of privacy, a credible link to an institution of higher education, reference to its standards of research ethics, and no incentive for falsification, we think it reasonable to assume at least a loose relationship between subjects' answers and their actual information and beliefs. Without minimizing the problems in question, we shall proceed accordingly.³⁰

Of our subjects who completed the experiment from unique IP addresses, only 12% (by self-report) are full-time students age 22 or below, the majority of these being

²⁹ To further confirm this, we estimated an OLS regression, with first-mover earnings as dependent variable, the independent variables being first-mover sending and its square, treatment dummies, treatment dummies interacted with first-mover sending, and treatment dummies interacted with the square of first-mover sending. Whereas the coefficients on first-mover sending and its square are both positive and jointly significant (the coefficient values are 0.401 and 0.006, respectively), no corresponding pair of interacted terms (for example, COOP*send and COOP*send²) is jointly significant. Results are shown in Appendix Table A.2.

³⁰ Ways of verifying subjects' identities through a background check can be devised, but since these are time-consuming, they would discourage participation without offsetting compensation. Also, because it would be difficult to verify characteristics like age and gender without the risk of revealing personal identity, such verification would also conflict with anonymity, which is a central element of the trust game design (since the experimenter does not want subjects "acting nice" for the sake of reputation). Note that the survey questions, which are shown in the Appendix, distinguish clearly between questions about the characteristics of the subjects' "in world" identities, i.e. their avatars, and questions about their "first life" identities or real world selves. We compare our subjects' self-reported demographics with those of Second Life residents in general in footnote 32.

American. 29.5% of subjects report real ages of 23 – 30 years, 24.6% 31 – 40, 12.6% 41 – 50, and 6.5% 51 or older. 15% of subjects report having a post-college degree, 30% being college graduates, and another 46% high school graduates. 17% report being full-time students, 45% report being in full-time employment, 20.7% part-time students and/or in part-time employment, and 17% report being unemployed and not in school.

In terms of real-world incomes, employed subjects reported a median hourly wage of \$12-\$15 an hour for real-world employment, with 20% earning under \$5.85 an hour and 29% more than \$20 an hour.³¹ Assuming a modicum of accuracy in these responses, this suggests that a considerable fraction of the subject pool could have been induced to report to a brick-and-mortar experiment lab only with a considerable cash outlay by the experimenter. 54% reported their real-world gender to be female, with 90% reporting that their avatar's and real-world gender were the same. About 59% of subjects had US-based IP addresses, 33% addresses in Europe, Canada or Australia, 4% addresses in Asian countries, and 3% ones in Latin America, with the remaining 1% from other countries or unidentified locations.³²

To help us understand subjects' valuation of Linden dollars for purposes of assessing the possibility that experimenters might leverage the purchasing power of their subject payment dollars by paying with "in world" currency, our post-task survey questions asked subjects about their sources of L\$'s. 55% of subjects reported being employed for L\$'s in Second Life, with 93% of the latter earning 0 to L\$500 per hour in such employment, 5% L\$501 to L\$2,000 an hour, and 2% over L\$2,001 an hour. Of those employed in Second Life, 48% reported doing less than 5 hours of such paid work

³¹ Of the remainder, 10% reported earnings between \$5.85 and \$8, 17% between \$8 and \$12, 11% between \$12 and \$15, and 12% \$15 to \$20. Of those earning more than \$20 an hour, 18% reported earning \$20 to \$40 an hour, and 11% over \$40/hr.

³² One way to roughly check the accuracy of our subjects' responses is to compare their self-reported demographics to those reported independently for Second Life residents as a whole. Relying on various sources, including Linden Labs, Borst, 2009, reports a distribution of hours spent in Second Life by country of residence as 39.4% for the United States, 41.9% for Europe, Canada or Australia, and 5.5% for Japan. The greater bias of our sample toward the U.S. may partly be explained by the language factor (Borst reports that some 34% of SL residents had language settings of continental European languages, while 3% had language settings of Japanese). The age distribution reported uses different ranges than those in our survey but shows a strong resemblance to our numbers: less than 1% ages 13 – 17, 15.1% ages 18 – 24, 34.5% ages 25 – 34, 28.5% ages 35 – 44, and 21.1% age 45 and over. The reported gender breakdown is somewhat heavier on males than our sample's self-reports (58.7% male, 41.3% female) but the difference is not dramatic. Overall, these and other independent reports of demographics appear similar enough to those obtained by us, especially if we keep in mind that our participants were self-selected out of the broader SL population.

per week, 36% 6 to 20 hours, and 16% 21 hours or more per week. 27% said they had purchased L\$'s for a real-world currency during the previous month. For the regression exercises to which we turn next, we use the responses on hours worked and currency purchased to create an indicator we call "in world money" which equals one for subjects who reported having a job in Second Life and not purchasing Lindens with real world currency during the previous month. 40% of participants (41% of the first movers and 39% of the second movers) meet this criterion.

Does it matter whether a subject is a student or an older, perhaps employed, individual? We begin to explore the influence of student or non-student status,³³ age, gender, real-world earnings, and source of L\$'s upon decisions in the experiment by including these factors as explanatory variables, along with treatment dummies, in Tobit regressions in which the dependent variable is the amount of endowment sent by a first-mover (Table 2) or the proportion of the amount received that is returned by a second-mover (Table 3). The second-mover regressions include only subjects who received positive amounts, and the amount received is included as an additional control. When only a single personal characteristic is entered alongside the treatment dummies, self-reported gender has no effect, while student status attains a highly significant and economically large negative coefficient for both first-mover sending and second-mover fraction returned, implying that students were significantly less trusting and trustworthy (less generous) in their decisions. Perhaps for parallel reasons, chronological age obtains positive coefficients in the two regressions in which it enters alone among the individual characteristics (although its coefficient falls short of significance in the regression for first-mover sending), and self-reported real-world income is also a significant positive predictor of amount sent and proportion returned. Our indicator that the subject is likely to have obtained his or her virtual currency by working in Second Life rather than by purchase on the LindeX™ for U.S. dollars, "in-world money," obtains a positive coefficient falling just short of the 10% significance level in Table 2 and a negative coefficient significant at the 10% level in Table 3, when entered singly. When all five

³³ For purposes of this exercise, the category "student" includes not only the 12% of subjects reporting their age to be 18-22 and their occupational status to be "full-time student" but also the 8% of subjects reporting their age to be 18 – 22 and their occupational status "part-time student and/or part-time work" as well as the 1% of subjects who report in the post-decision survey an age below 18 despite declaring their age to be 18 or above at the participant consent stage before the experiment began.

individual-level variables are entered simultaneously, as shown in column (7) of each table, the coefficient on “in-world money” becomes significant and retains its opposite signs in both tables, while the student, age and income variables retain their signs in all but one case, with at least one of them (student, for the first-mover regression, age, for the second-mover one) being individually significant.

In addition to studying the general effects of subject characteristics on their experimental trusting and trustworthiness, our data also allow us to investigate whether subjects having different demographic characteristics responded differently to our visual and textual treatments. One approach is to include treatment-characteristic interaction terms (for example, Female*COOPERATIVE) in regressions like those of tables 2 and 3, but the number of potential interactions is large relative to the number of observations and no significant coefficients emerged in this exercise. As an alternative, we split the sample according to each criterion (male versus female, student versus non-student, etc.), making the sub-groups about equally sized when possible,³⁴ and performed Chow tests of whether the relationships captured by the regressions specified in the last columns of each table can be treated as the same in each paired sub-population. Three pairs of regressions were found to differ between sub-populations: the first-mover regression differs between male and female subjects ($p = .0481$); and both first-mover and second-mover regressions differ between younger and older subjects ($p = .0999$ and $p = .0288$, respectively). For these three cases, Table 4 reports the regressions by sub-group.

Viewing the results, we find that even in these cases in which there are some differences between sub-populations, most of the main results hold for both sub-groups—e.g., there are qualitatively similar differences in first-mover sending between COOPERATIVE and COMPETITIVE treatments and also between OPTIMISTIC and

³⁴ For the age and income variables, we used the median observation as point of division and placed those at the median into whichever group (above, or below, median) would make the resulting sub-groups of more nearly equal size. Accordingly, participants with the median age (26.5 for first and second movers) were included with the younger split, with 88 first-movers in the younger group, 68 first-movers in the older group, 78 second-movers in the younger group and 64 second-movers in the older group. The median income was \$13,500, and this group was included with the lower earners for both first and second movers ($n=71$, and 74, respectively with 85 and 68 in the higher income group). Among first-movers, there were 82 males and 74 females, and 92 subjects who met the criterion for “in-world money” while 64 did not. Among second-movers, there were 55 males and 87 females, and 89 meeting the “in-world money” criterion versus 53 not doing so. Only the student/ non-student split unavoidably yielded groups of highly unequal size: 32 students and 124 non-students among the first-movers, and 31 students and 111 non-students among the second-movers who had received positive amounts.

CAUTIONARY treatments in both the male and the female and in both the younger and older sub-populations. The take-home message from this may be that it is not critical to reach beyond the typical subject pool of college-aged students, but we do not view this as a “no news” result since it is not something that could have been presumed without investigation. There are, moreover, some differences between groups.

The third column of each set shows the p -value of a test for differences in coefficients. These test results indicate that female first-movers sent significantly more under the COOP-OPTIMISTIC treatment than did males. Indeed, comparison of the treatment dummies for the COOPERATIVE, OPTIMISTIC and COOP-OPTIMISTIC treatments suggests that positive visual and textual stimuli had reinforcing effects on first-mover sending among female subjects, though that was not the case for male subjects nor for the sample as a whole. The negative effect of student status on first-mover sending is also significantly larger for female than for male subjects. For first-mover sending split by age, only the general effect of being female appears to differ significantly, being positive for older and negative for younger subjects—though this gender effect is not independently significant in the regression for either sub-population.

With regard to the proportion returned by second-movers, the regressions for which differ only under the division by age group, there is a significant difference in the coefficient on the CAUTIONARY treatment, which is more negative for younger subjects, and on the general effect of income, which is significantly negative for the younger sub-group and positive but insignificant for the older subjects. We also checked whether the *differences* between paired coefficients significantly differed between sub-groups. Focusing on the core treatment differences, those between the COOPERATIVE and COMPETITIVE treatments and those between the OPTIMISTIC and CAUTIONARY treatments, we found only one case of significant difference: the difference between the OPTIMISTIC and the CAUTIONARY treatment dummies is larger for younger than for older second-movers, significant at the 1% level, according to an F-test. While only one of the individual treatment dummy variables concerned is itself statistically significant, and that one only marginally so, this test indicates that, relative to older second-movers, younger second-movers were moved towards (away from)

reciprocity by the presence of the panda (fish) photo, while if anything the opposite was the case for older second-movers.³⁵

Survey trust instruments and trust game decisions

A large empirical literature on trust and macroeconomic outcomes has grown up independently of the experimental trust game literature, and several scholars, including Glaeser *et al.* (2000) and Sapienza *et al.* (2007), have investigated how, if at all, experimental behaviors correlate with the survey responses. In our post-task survey questions, we included the generalized trust question from the World Values Survey (“Generally speaking, would you say that most people can be trusted or that you can’t be too careful in dealing with people?”) and we also asked our subjects whether, in their opinion, “Second Life residents [are] more trustworthy, less trustworthy, or about the same as people in the general population?” While the responses are uncorrelated with first-mover choices, they are highly correlated with the proportion of received money returned by second-movers: subjects who selected the trusting response to the generalized trust question and subjects who rated Second Life residents as more trustworthy were themselves more trustworthy in their decisions, mirroring closely the finding of Glaeser *et al.* (2000) that the trust question predicts trustworthiness but not trust.³⁶ In contrast, answers to the question of whether Second Life residents are “more worthy of reciprocal

³⁵ An F-test of the right-most regression in Table 4 shows that although the relevant coefficients are individually insignificant, the coefficients on the OPTIMISTIC and CAUTIONARY treatment dummies differ from each other at the 5% significance level for older subjects, with such subjects actually tending to return a higher proportion of the amount received under the CAUTIONARY than under the OPTIMISTIC treatment. It’s worth bearing in mind that OPTIMISTIC treatment subjects had received considerably more from their first movers than had their CAUTIONARY treatment counterparts and that returning a smaller proportion does not imply returning a smaller amount. Perhaps some older second movers were more disposed to reward sending that took place in a cool and discouraging environment than were their counterparts disposed to reward sending in a warm and cuddly one because the former perceived sending in the environment in question as a more impressive act of trust than did the latter.

³⁶ We estimated tobit regressions with either sending by A or proportion returned by B as dependent variable, with the independent variables being the responses to the generalized trust and other opinion questions entered either jointly or singly, plus treatment dummy variables. In one set of regressions for second-mover behaviors, we included also the amount received from A. The other opinion variables concern the trustworthiness of Second Life residents, and the deservingness of Second Life residents, in both cases relative to the general population. In the proportion returned but not in the amount sent regressions, the first two opinion variables obtained statistically significant coefficients whether entered jointly or singly. Results are shown in Appendix Tables A.3.1 and A.3.2. We note that the responses to the opinion questions by second-movers are uncorrelated with the amounts they had received from their counterparts (already known to them when answering), so stated views on trust whether in general or among Second Life residents appear to be unaffected by the counterpart’s choice.

treatment ... than people in the general population?” were uncorrelated with either first- or second-movers’ behaviors.

Discussion

We investigated the potential to conduct economic experiments in a virtual world, Second Life, choosing as our example the trust game introduced by Berg *et al.* (1995). One motivation was to consider the possibility that a more diverse subject pool could be recruited and engaged in an experiment with a high degree of anonymity and with effective incentives at reasonable cost to the experimenter. While we forewent verification of subject characteristics for reasons of cost and privacy, it appears that we succeeded in obtaining a subject population that is considerably more diverse than that in the usual laboratory experiment in terms of age, occupation, and education. According to our data, age and employment status do make a difference to subjects’ general behaviors, which confirms the desirability of a broad subject pool. But the differences run contrary to suspicions that students are unrepresentatively naïve or unselfish. Despite an unusually high degree of anonymity, most subjects did not adopt “asocial” patterns of behavior. On the contrary, there was as much trusting as in lab and field trials of the same experiment, and somewhat more trustworthiness.

We used convenient features of the virtual lab environment to investigate subjects’ sensitivities to contrasting visual and textual modifications of the trust game. While replicating previous findings that the overwhelming majority of first-movers send money to their counterparts, we found their decisions to be sensitive to small differences in instruction text or viewed environment. The average amount sent by first-movers was almost identical to that in BDM, but at the treatment level it differed by about a quarter of the endowment (or about half of the average amount sent) as between two visual cue treatments selected to encourage an optimistic social outlook or a cautious self-interested outlook, respectively. Average sending differed by a similar amount between two add-on text treatments, one suggesting possibilities of cooperation, the other spelling out the implications of individual payoff maximization.

Second-movers were exposed to the same treatment differences as their first-mover counterparts but the amounts they returned—which sufficed to make trusting pay

off on average—varied significantly mainly as a function of amount received, not treatment. Apparently, the norm of reciprocating the trust shown by a counterpart was much stronger than any impact of the verbal and visual cues distinguishing our treatments. And while first-movers faced no stimulus to action other than the instructions and visual environments, small manipulations of which proved quite influential, second-movers had the additional stimulus of an actual counterpart's decision, which evidently was far more salient to them than were our changes in cues.

In addition to the general findings just noted, we also found some differences in treatment effects between male and female and between younger and older subjects. Positive visual and textual stimuli were reinforcing when combined for female but not male first-movers. Positive and negative visual stimuli triggered different effects on the proportion returned by younger than by older second movers

Like Bateson *et al.* (2006), our visual treatment comparisons strongly support the idea, suggested by evolutionary psychology, that human social behaviors are responsive to environmental cues. This finding fits well with the interpretation of trust game behaviors as being much influenced by social predispositions, and we interpret the replication of such behaviors even in our highly anonymous design in a virtual world setting as a further indication of those dispositions' strength. We found that verbal suggestions could have an equally large influence on behavior as did visual ones, but the two kinds of influence operated asymmetrically. The fact that the visual cues were more successful at bolstering sociality while the verbal ones were better at quashing it suggests that different pathways are at work, with the visual cue best at strengthening partly unconscious social intuitions and the textual one best at activating conscious rationality to over-ride those intuitions. Of course, our small sample and the specific nature of the visual and textual cues used in our experiment must be borne in mind, so any extrapolation from the conclusions just stated should be done with great caution.

One issue likely to be of interest for future research is whether we were successful in providing meaningful economic incentives to our subjects despite the lower dollar value of the incentives provided. A common yardstick with which to measure the scale of payoffs is the comparison of subjects' earnings to average wage rates in the environment from which they are drawn. About half of our subjects reported holding paid jobs in

Second Life, and their average hourly pay rates in those jobs are about a third of the average earnings rate from our experiment, converted to an hourly basis. That our subjects' trust game behaviors are not dramatically different from those in BDM or in the contemporaneous Second Life trust game experiments of Fiedler and collaborators and of Chesney *et al.*, with considerably higher stakes, may be taken as another indication of incentive adequacy. More pertinent, perhaps, is the fact that the large majority of subjects reported that the prospect of earnings was at least a moderately important motivation for participation, that they had good knowledge of the range of possible earnings when they entered the experiment, and that they found the range of earnings to be reasonable. While we remain cautious about drawing strong conclusions from these findings, there appears to be enough evidence of incentive adequacy to justify exploration by means of further experimentation, which might involve both variation of earnings ranges and more intensive debriefing of subjects.

References

- Arrow, Kenneth, 1974, *The Limits of Organization*. New York: Norton.
- Atlas, Stephen A., 2008a, "Advancing Toward Inductive Metanomics: Economic Experiments on Framing, Reciprocity and Trust in Virtual Worlds," Master Thesis, Tufts University Department of Economics.
- Atlas, Stephen A., 2008b, "Inductive Metanomics: Economic Experiments in Virtual Worlds," *Journal of Virtual Worlds Research*, 1 (1): 1 – 15.
- Bargh, John and Tanya Chartrand, "The Unbearable Automaticity of Being," *American Psychologist* 54 (7): 462-479.
- Barr, Abigail, 2003, "Trust and expected trustworthiness: Experimental evidence from Zimbabwe," *Economic Journal* 113, 614-630.
- Bateson, Melissa, Daniel Nettle and Gilbert Roberts, 2006, "Cues of Being Watched Enhance Cooperation in a Real-World Setting," *Biological Letters* 2: 412-414.
- Benbasat, Izak, David Gefen and Paula Pavlou, 2008, "Special Issue: Trust in Online Environments," *Journal of Management Information Systems* 24 (4): 5 – 11.
- Ben-Ner, Avner and Louis Putterman, 1998, "Values and Institutions in Economic Analysis," in Ben-Ner and Putterman, Eds., *Economics, Values and Organization*. New York: Cambridge University Press.
- Ben-Ner, Avner, Louis Putterman and Ting Ren, 2007, "Lavish Returns on Cheap Talk," Working Paper 2007-15, Department of Economics, Brown University.
- Ben-Ner, Avner and Louis Putterman, 2009, "Trust, Communication and Contracts," *Journal of Economic Behavior and Organization* 70 (1-2): 106-121.
- Berg, Joyce, John Dickhaut, and Kevin McCabe, 1995, "Trust, Reciprocity and Social History," *Games and Economic Behavior* 10: 122-42.
- Bloomfield, Robert and Kristina Rennekamp, 2009, "Experimental Research in Financial Reporting: From the Laboratory to the Virtual World," *Foundations and Trends in Accounting*. 3 (1): 1-85.
- Borst, Todd, 2009, "Know Your Customers: Second Life Demographics," Virtual World Business.Com, April 30, accessed at <http://virtualworldbusiness.com/2009/04/30/know-your-customers-second-life-demographics/>
- Boyd, Robert, and Peter Richerson, 1985, *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.

Brühlhart, Marius and Jean-Claude Usunier, “Verified Trust: Reciprocity, Altruism and Noise in Trust Games,” CEPR Discussion Paper No. 4758, Center for Economic Policy Research.

Burks, Stephen, Jeffrey Carpenter and Eric Verhoogen, 2003, “Playing Both Roles in the Trust Game,” *Journal of Economic Behavior and Organization* 51 (2): 195-216.

Burks, Stephen, Jeffrey Carpenter, Lorenz Goette, Kristen Monaco, Kay Porter, and Aldo Rustichini, 2007 “Using Behavioral Economic Experiments at a Large Motor Carrier: The Context and Design of the Truckers and Turnover Project,” IZA Discussion Paper No. 2789.

Burnham, Terence, Kevin McCabe and Vernon Smith, 2000, “Friend-or-foe Intentionality Priming in an Extensive Form Trust Game,” *Journal of Economic Behavior and Organization* 43: 57-73.

Buss, David, 2008, *Evolutionary Psychology. The New Science of the Mind, 3rd Edition*. Boston: Allyn and Bacon.

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

Carter, Michael and Marco Castillo, 2002, “The Economic Impacts of Altruism, Trust and Reciprocity: An Experimental Approach to Social Capital,” University of Wisconsin-Madison Agricultural and Applied Economics Staff Paper.

Castronova, Edward, 2008, “A Test of the Law of Demand in a Virtual World: Exploring the Petri Dish Approach to Social Science,” *CESifo Working Paper Series No. 2355*. Available at SSRN: <http://ssrn.com/abstract=1173642>

Chesney, Thomas, Swee-Hoon Chuah and Robert Hoffmann, 2007, “Virtual World Experimentation: An Exploratory Study,” Working Paper, Available at SSRN: <http://ssrn.com/abstract=1068225>

Chesney, Thomas, Swee-Hoon Chuah and Robert Hoffmann, 2009, “Trust in V-Commerce: An Experimental Approach,” mimeo, International Center for Behavioral Business Research, Nottingham University Business School.

Cox, James, 2004, “How to identify trust and reciprocity,” *Games and Economic Behavior* 46: 260-81.

Dal Bó, Ernesto and Pedro Dal Bó, 2009. “Do the Right Thing: The Effect of Moral Suasion on Cooperation,” unpublished paper, Haas School of Business.

- Danielson, Anders and Hakan Holm, 2005, "Tropic Trust Versus Nordic Trust: Experimental Evidence from Tanzania and Sweden," *Economic Journal* 115 (503): 505-532.
- Eckel, Catherine and Rick Wilson, 2006, "Internet Cautions: Experimental Games with Internet Partners," *Experimental Economics* 9: 53-66.
- Fehr, Ernst and John List, 2004, "The Hidden Costs and Returns of Incentives—Trust and Trustworthiness among CEOs," *Journal of the European Economic Association* 2(5): 743-71.
- Fiedler, Marina, 2009, "Cooperation in Virtual Worlds," *Schmalenbach Business Review* 61: 173-194.
- Fiedler, Marina and Ernan Haruvy, 2009, "The Lab versus the Virtual Lab and Virtual Field—An Experimental Investigation of Trust Games with Communication," *Journal of Economic Behavior and Organization* 72 (2): 716-724.
- Fiedler, Marina, Ernan Haruvy and Sherry Xin Li, 2008, "Social Distance in a Virtual World Experiment," unpublished paper, Ludwig-Maximilians-University and University of Texas at Dallas.
- Foster, Andrea, 2007, "Virtual Worlds as Social Science Labs," *Chronicle of Higher Education*. 6 July 2007. 53.44. A25.
- Glaeser, Edward, David Laibson, José Scheinkman and Christine Soutter, 2000, "Measuring Trust," *The Quarterly Journal of Economics* 65, 811-846.
- Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, and Richard McElreath, 2001, "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies," *American Economic Review* 91(2): 73-78.
- Herrmann, Benedikt, Christian Thöni and Simon Gächter, 2008, "Antisocial Punishment Across Societies," *Science* 319: 1362-7.
- Hoffman, Elizabeth, Kevin McCabe and Vernon Smith, 1998, "Behavioral Foundations of Reciprocity: Experimental Economics and Evolutionary Psychology," *Economic Inquiry* 36: 335-52.
- Holm, Hakan and Anders Danielson, 2005, "Tropic Trust Versus Nordic Trust: Experimental Evidence from Tanzania and Sweden," *Economic Journal* 115 (503): 505-532.
- Karlan, Dean, 2005. "Using experimental economics to measure social capital and predict financial decisions," *American Economic Review* 95(5):1688-99.

Kirchsteiger, Georg, Luca Rigotti and Aldo Rustichini, 2006, "Your Morals Might be Your Moods," *Journal of Economic Behavior and Organization* 59: 155-172.

Klauer, Karl, 1997, "Affective Priming," *European Review of Social Psychology* 8 (1): 67-103.

Linden Lab, 2008, "Second Life Key Metrics through June 2008," Acc Dec 29, 2008: <http://spreadsheets.google.com/pub?key=pxbDc4B2FH94lKcWU2YDVFQ&gid=1>

Page, Talbot and Louis Putterman, 2010, "An Uninterpreted Spatial Version of the Trust Game: Evidence of Reciprocity without Suggestive Words, Evidence of Iterated Dominance Self-Taught," mimeo, Brown University.

Ross, Lee and Andrew Ward, 1996, "Naïve Realism in Everyday Life: Implications for Social Conflict and Misunderstanding," in Edward Reed, Elliot Turiel, and Terrance Brown, eds., *Values and Knowledge*. Mahwah, NJ: L. Erlbaum.

Sapienza, Paola, Anna Simats and Luigi Zingales, 2007, "Understanding Trust," CEPR Discussion Paper No. DP6462.

Scharlemann, Jörn, Catherine Eckel, Alex Kacelnik and Rick Wilson, 2001, "The Value of a Smile: Game Theory with a Human Face," *Journal of Economic Psychology* 22: 617-640.

Schechter, Laura, 2007, "Traditional Trust Measurement and the Risk Confound: An Experiment in Rural Paraguay," *Journal of Economic Behavior and Organization* 62 (2): 272-292.

Wright, Robert, 1994, *The Moral Animal: The New Science of Evolutionary Psychology*. New York: Vintage Books.



Figure 1. Picture of ‘lobby’ area of laboratory. Subjects first clicked on the automated greeter to be welcomed to the lab, then on the booth to the right to be transported to a private treatment room.

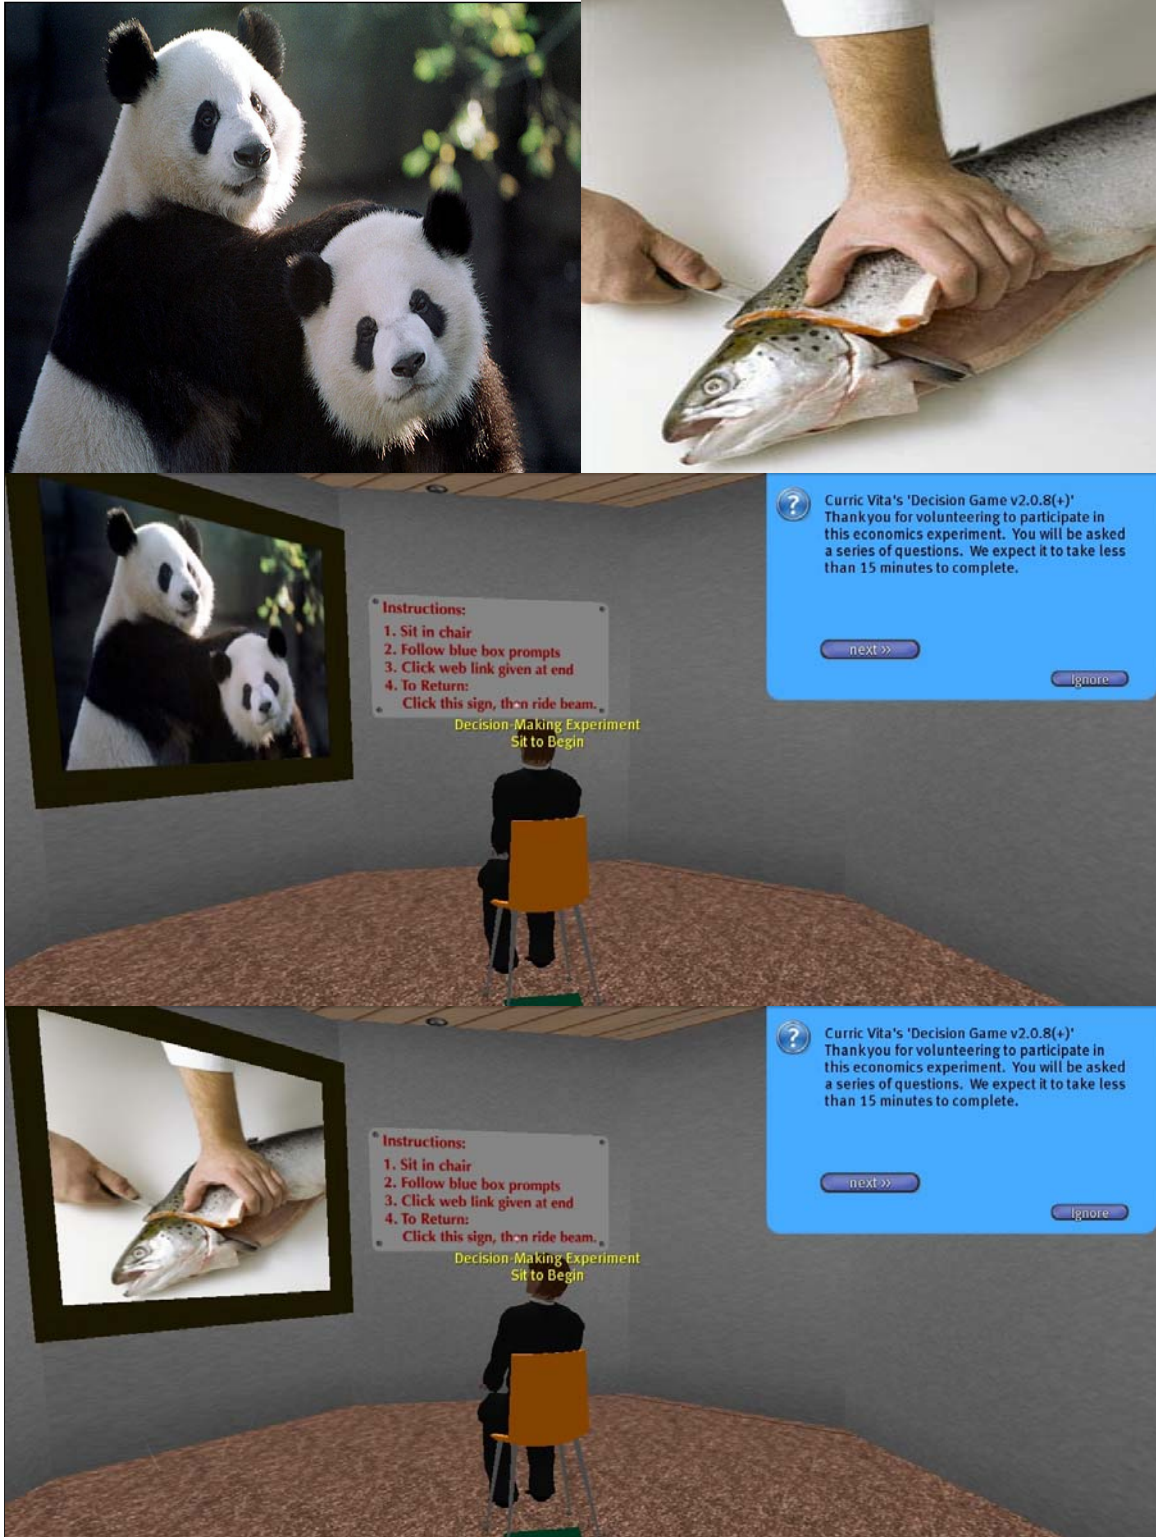


Figure 2. “OPTIMISTIC” and “CAUTIONARY” photos (top) and views of experiment rooms in corresponding treatments.

Table 1a. Number of non-duplicate-IP observations, mean and standard deviations of sending by A (X_A), for all observations and those with nonzero sending ($X_A > 0$), by treatment.

Treatment	No. of A Obs.		Mean and Standard Deviation of Amt. Sent			
	All	$X_A > 0$	Mean (All)	St. Dev.	Mean $X_A > 0$	St. Dev
BASELINE	19	19	49.3	33.9	49.3	33.9
COOPERATIVE	19	18	59.1	39.1	62.4	37.4
COMPETITIVE	21	17	33.1	29.3	40.1	27.2
OPTIMISTIC	23	22	70.4	36.0	73.6	33.3
CAUTIONARY	29	27	48.9	38.4	52.5	37.2
COOP- OPTIMISTIC	21	20	55.0	37.4	57.8	36.2
COMP- CAUTIONARY	24	23	47.0	27.5	49.0	26.1
Overall	156	146	51.8	35.7	55.3	34.2

Table 1b. Number of non-duplicate-IP observations, mean and standard deviations of sending by B (X_B), and B sending as a fraction of non-zero amount received ($X_B/3X_A$, given $X_A > 0$), by treatment

Treatment	No. of B Obs.		Mean Amt. Sent		$X_B/3X_A$ when $X_A > 0$	
	All	Received > 0	Mean	St. Dev.	Mean	St. Dev
BASELINE	23	23	86.3	80.9	.507	.305
COOPERATIVE	22	20	108.18	86.8	.570	.199
COMPETITIVE	16	12	41.8	56.0	.399	.323
OPTIMISTIC	20	19	87.1	85.5	.443	.330
CAUTIONARY	28	26	56.7	63.7	.418	.294
COOP- OPTIMISTIC	21	20	78.8	79.2	.464	.315
COMP- CAUTIONARY	23	22	74.2	66.3	.488	.307
Overall	153	142	76.6	76.0	.473	.296

Note: The number of B observations does not equal the number of A observations in some cases due to dropped observations from non-unique IP addresses.

Table 2. First-mover sending as a function of treatment and self-reported characteristics (Tobit regressions)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
COOP dummy	16.8 (15.9)	16.5 (15.9)	15.4 (15.4)	17.1 (15.7)	24.8 (15.1)	15.5 (15.8)	21.7 (14.9)
COMP dummy	-24.6 (15.2)	-24.9 (15.3)	-27.6* (14.9)	-23.8 (15.1)	-14.3 (14.4)	-22.1 (15.2)	-15.7 (14.3)
OPTIMISTIC dummy	33.7** (15.4)	33.8** (15.4)	35.1** (14.9)	36.2** (15.3)	46.9*** (14.8)	32.4** (15.3)	44.1*** (14.5)
CAUTIONARY dummy	-0.86 (14.2)	-1.0 (14.2)	-4.2 (13.8)	0.3 (14.0)	7.4 (13.4)	-0.6 (14.1)	3.4 (13.3)
COOP-OPTIMISTIC dummy	9.7 (15.3)	9.8 (15.3)	8.0 (14.9)	9.0 (15.2)	23.5 (14.7)	11.1 (15.2)	23.2 (14.6)
COMP-CAUTIONARY dummy	-5.2 (14.6)	-5.4 (14.6)	-7.6 (14.3)	-4.2 (14.5)	3.9 (13.9)	-5.0 (14.5)	0.7 (13.7)
Female		2.6 (7.9)					1.0 (7.3)
Student			-25.0*** (9.5)				-21.0* (11.1)
Age				0.6 (0.4)			-0.5 (0.4)
Income					0.9*** (0.2)		0.8*** (0.2)
“in-world money”						13.3 (8.1)	12.8* (7.5)
No. of obs.	156	156	156	156	156	156	156
Log likelihood	-612.1	-612.0	-608.7	-610.9	-603.6	-610.7	-600.5
Prob > χ^2	0.0073	0.0134	0.0010	0.0058	0.0000	0.0050	0.0000
Pseudo R ²	0.0142	0.0143	0.0196	0.0160	0.0278	0.0163	0.0328

Notes: All regressions have 156 observations. Omitted treatment dummy is BASELINE. Numbers in parentheses are standard errors. *** = significant at 1% level; ** = significant at 5% level; * = significant at 10% level.

Table 3. Second-mover proportion returned as a function of treatment and self-reported characteristics (Tobit regressions)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Fraction sent by first-mover	0.017 (0.095)	0.020 (0.097)	0.002 (0.093)	0.001 (0.092)	-0.007 (0.093)	0.045 (0.095)	0.021 (0.092)
COOPERATIVE dummy	0.084 (0.112)	0.083 (0.112)	0.073 (0.109)	0.057 (0.108)	0.071 (0.109)	0.088 (0.110)	0.056 (0.105)
COMPETITIVE dummy	-0.111 (0.131)	-0.110 (0.131)	-0.104 (0.128)	-0.142 (0.126)	-0.142 (0.128)	-0.131 (0.130)	-0.170 (0.124)
OPTIMISTIC dummy	-0.055 (0.116)	-0.057 (0.116)	-0.052 (0.113)	-0.032 (0.111)	-0.065 (0.113)	-0.051 (0.114)	-0.043 (0.109)
CAUTIONARY dummy	-0.090 (0.105)	-0.091 (0.105)	-0.073 (0.103)	-0.081 (0.101)	-0.133 (0.104)	-0.072 (0.104)	-0.091 (0.101)
COOP-OPTIMISTIC dummy	-0.030 (0.113)	-0.030 (0.112)	-0.033 (0.110)	-0.035 (0.108)	0.056 (0.110)	-0.030 (0.111)	-0.051 (0.101)
COMP-CAUTIONARY dummy	-0.012 (0.109)	-0.013 (0.109)	-0.020 (0.107)	-0.032 (0.105)	-0.049 (0.108)	-0.015 (0.108)	-0.054 (0.103)
Female		-0.012 (0.065)					-0.019 (0.061)
Student			-0.174** (0.074)				-0.043 (0.087)
Age				0.010*** (0.003)			0.006* (0.004)
Income					0.005** (0.002)		0.003 (0.002)
“in-world money”						-0.121* (0.065)	-0.122** (0.062)
No. of obs.	142	142	142	142	142	142	142
Log likelihood	-81.54	-81.52	-78.81	-76.36	-78.26	-79.83	-73.27
Prob > χ^2	0.8248	0.8886	0.3373	0.0828	0.2533	0.5341	0.0645
Pseudo R ²	0.0216	0.0218	0.0544	0.0837	0.0610	0.0421	0.1208

Notes: All regressions have 142 observations, that is all complete observations except those in which the first-mover sent zero. Omitted treatment dummy is BASELINE. Numbers in parentheses are standard errors. *** = significant at 1% level; ** = significant at 5% level; * = significant at 10% level.

Table 4. First-mover sending and second-mover proportion returned as functions of treatment and self-reported characteristics, for gender and age-segmented sub-samples (Tobit regressions)

	First Mover Sending, Segmented by Gender			First Mover Sending, Segmented by Age			Second Mover Sending, Segmented by Age		
	Male	Female	<i>p</i> -value	Younger	Older	<i>p</i> -value	Younger	Older	<i>p</i> -value
Cooperative	21.7 (20.6)	23.6 (19.9)	0.945	6.7 (19.6)	33.6 (22.6)	0.3709	0.07 (0.14)	0.06 (0.14)	0.9797
Competitive	-17.8 (19.9)	-16.1 (18.8)	0.951	-18.9 (19.7)	-14.7 (19.3)	0.8791	-0.02 (0.19)	-0.18 (0.15)	0.5165
Optimistic	55.3 (18.4)***	25.0 (20.6)	0.273	54.9 (18.3)***	26.2 (23.1)	0.3304	0.04 (0.13)	-0.20 (0.17)	0.2666
Cautionary	11.2 (18.0)	-1.9 (17.9)	0.607	-1.9 (17.7)	0.6 (19.2)	0.9246	-0.23 (0.12)*	0.16 (0.15)	0.0461**
Coop-Optimistic	5.0 (19.0)	57.3 (21.1)***	0.068*	16.1 (19.5)	32.3 (20.2)	0.5647	-0.19 (0.13)	0.11 (0.16)	0.1513
Comp-Cautionary	9.9 (18.4)	-12.6 (18.6)	0.389	8.8 (18.2)	-20.5 (19.4)	0.2730	-0.11 (0.13)	0.02 (0.15)	0.5288
First-mover percent sent							0.01 (0.12)	0.08 (0.14)	0.7052
Female				-10.5 (9.6)	16.6 (11.1)	0.0665	0.05 (0.08)	-0.08 (0.09)	0.2831
Age	-0.2 (0.6)	-0.4 (0.6)	0.841						
Student	-4.9 (13.9)	-50.0 (16.6)	0.039**	-19.2 (10.3)*	(omitted)		-0.14 (0.08)*	(omitted)	
Income	0.7 (0.3)**	1.0 (0.3)***	0.499	1.5 (0.4)***	0.8 (0.3)***	0.1651	-0.00 (0.004)	0.01 (0.00)**	0.0308**
IGL	18.3 (9.7)*	5.2 (10.6)	0.364	21.5 (9.6)**	4.2 (12.2)	0.2670	-0.12 (0.08)	-0.17 (0.09)*	0.6385
Intercept	29.6 (25.9)	46.7 (26.6)*	0.648	25.7 (18.0)	20.0 (19.3)	0.8303	0.59 (0.13)***	0.43 (0.14)***	0.4090
No. of obs.	82	74		88	68		78	64	
Log-likelihood	-590.40			-592.72			-63.68		
Chow test <i>p</i> >F	0.0481**			0.0999 *			0.0288**		

Numbers in parentheses are standard errors. *** = significant at 1% level; ** = significant at 5% level; * = significant at 10% level. Median age was 26.5, with 48 first-movers and 39 second-movers in this category. These individuals were grouped with the “younger” group to minimize differences in segment size. Reported *p*-values are those for F-tests of differences in coefficients across sub-sample regressions.