

# Who do you distrust and how much does it cost? An experiment on the measurement of trust

**Roberto A. Weber**

*Department of Social & Decision Sciences  
Carnegie Mellon University*

**Bill McEvily**

*Rotman School of Management  
University of Toronto*

**Joseph R. Radzevick**

*Tepper School of Business  
Carnegie Mellon University*

July 9, 2008\*

We address two problems with how trust is frequently measured in economics. First, we highlight the importance of clearly identifying the target of trust, which when ignored can lead to inconsistencies between trust measures. Second, we note the importance of distinguishing trust from other closely related concepts. We conduct an experiment using a new behavioral measure of trust – individuals' willingness to pay to avoid being vulnerable to the target of trust – and vary the target of trust. To test our behavioral measure, we also collect data on potentially confounding effects (i.e., altruism and risk aversion) and on attitudinal measures of trust. Subjects discriminate based on perceived characteristics of different targets in determining whether to trust, in a manner consistent with trust elicited using attitudinal measures and with actual trustworthiness. Risk aversion and altruism do not correlate highly with our measure of trust.

---

\* We gratefully acknowledge funding from a Carnegie Mellon Berkman Faculty Development Grant. Weber also gratefully acknowledges support from the National Science Foundation (SES-0433152). We also thank participants at several seminars and conferences for helpful comments and suggestions. Finally, we are also very thankful to John Hamman for valuable research assistance, the Pittsburgh Experimental Economics Laboratory (PEEL) for access to resources, and David Laibson for granting us access to his data. Finally, one of the authors would like to thank his wife, Stephanie, for valuable support and help during the completion of this project.

## **I. Introduction**

Trust is often regarded as an important concept for understanding social, economic, organizational, and financial activity (Arrow 1972; Putnam 1993; Knack & Keefer 1997; McEvily et al 2003; Guiso et al 2004). Trust is viewed as a mechanism whereby potentially beneficial exchange can occur while overcoming the presence of moral hazard.

In spite of the potential importance of trust, there is little agreement across the social sciences on how it should be defined or, especially, on how it should be measured (Glaeser et al 2000; McEvily & Tortoriello 2007). For instance, the measurement of trust ranges from behavioral measures such as the “investment game” (Berg et al 1995), to attitudinal questions eliciting general perceptions of trustworthiness (Knack & Keefer 1997; La Porta et al 1997; Rotter 1967), to multi-item attitudinal survey measures eliciting trust in specific contexts and towards specific others (Cook & Wall 1980; Rempel & Holmes 1986; Currall & Judge 1995).

Whereas behavioral measures tend to be favored by economists, attitudinal measures are used most frequently in other social sciences (e.g., social psychology, sociology, and organizational research). This discrepancy in the methods used for measuring trust may not be a problem by itself, provided the different methods all measure the same concept and that concept is something researchers are comfortable interpreting as trust. However, if the various measures of trust are unrelated to each other or if they measure things other than trust, then social scientists using these measures must question the extent to which they are actually measuring trust.

*Do different measures of trust capture the same thing?* A few studies simultaneously explore behavioral and attitudinal methods of measuring trust. In general, these studies find little

relationship between attitudinal measures and behavioral measures such as the investment game (Glaeser et al 2000; Karlan 2005; Danielson & Holm 2007).<sup>1</sup>

In the most familiar of these studies, Glaeser et al (2000) tested the relationship between several behavioral and attitudinal measures and found very little reliability across different kinds of measures. Based on these findings Glaeser et al concluded that “. . . standard survey questions about trust do not appear to measure trust . . . This means that most work using these survey questions needs to be somewhat reinterpreted” (p. 841). Although Glaeser et al suggest the inferiority of attitudinal measures, their results could equally be interpreted as calling into question the accuracy of the behavioral measures used in their study. Regardless of which view is correct, the Glaeser et al study and others like it certainly raise concerns about the extent to which researchers using different methods of measuring trust are capturing the same concept.

*Do measures of trust capture something other than trust?* Aside from concerns about reliability across measurement methods, individual measures of trust also may suffer from the criticism that they actually measure something different from trust. For instance, behavior in the most frequently used behavioral trust measure, the investment game, can be interpreted as reflecting a preference for altruism or fairness. Thus, measuring trust using the investment game may involve disentangling motives for observed behavior (Cox 2004; Ashraf et al 2006; Capra et al 2007). Moreover, “trusting” behavior in the investment game may correlate significantly with risk seeking (Karlan 2005; Schechter 2007).<sup>2</sup>

---

<sup>1</sup> We know of two notable exceptions. Holm & Danielson (2005) find a positive relationship between attitudinal and behavioral trust for one population (Swedes), but not for another (Tanzanians), though it is unclear precisely why they find such a difference. Capra et al (2007) find that investment game behavior is correlated with attitudinal trust measures, but only when one controls for altruism.

<sup>2</sup> However, Eckel and Wilson (2004) find only a weak relationship between risk aversion and behavior in the investment game.

## **II. Revisiting the measurement of trust**

As the above discussion reveals, there are potentially significant problems with how trust is measured. Behavior in the investment game, the principal behavioral measure, fails to correlate with other (attitudinal) measures of trust. This alone may not be problematic if one accepts Glaeser et al's (2000) interpretation of the superiority of behavioral measures. When one also considers, however, that investment game behavior correlates with altruism and risk seeking, it raises the question of whether this measure really captures something we can accurately refer to as trust.

In this paper, we aim to make two advances to the way that trust is measured. First, we emphasize the need to clearly identify the target, or object, of trust, which may account for the apparent lack of reliability across trust measures found by Glaeser et al and others. As we discuss below, holding constant the target of trust in Glaeser et al's data suggests a relationship between attitudinal and behavioral measures. Second, we stress the need to devise better behavioral methods for measuring trust that are less strongly confounded with other closely related concepts such as altruism and risk seeking.

Our study measures trust in a new way, with these objectives in mind. To clearly specify the target of trust, we develop and test an experimental procedure in which laboratory subjects are each matched with several field participants at different locations in the same city. This allows us to measure how trust changes as the target changes, and the extent to which target-specific attitudinal and behavioral measures are related. To clearly identify trust, separately from concepts such as risk seeking and altruism, we develop a game based on the costs that individuals are willing to incur to mitigate their vulnerability to the target. We find in our experiment that subjects' trust varies depending on the target, that such variation is consistent with the actual trustworthiness of the

target, that our target-specific behavioral measure correlates with attitudinal trust measures, and that our behavioral measure is only weakly related to other concepts such as risk and altruism.

#### *A. The importance of the target of trust*

Sociologists (Blau 1964; Luhmann 1979; Lewis & Weigart 1985) and social psychologists (Johnson-George & Swap 1982; Rempel & Holmes 1986; Lewicki & Bunker 1996) have long noted the importance of the target, or object, of trust and the ability of individuals to discriminate the trustworthiness of different targets. For instance, individuals may exhibit high trust towards their family members but not towards the government; an employee may fully trust her direct supervisor but not the company CEO. From this perspective, trust is less a stable attribute, or disposition, of the person placing trust and more a property of each specific trustor-trustee interaction. Given this, sociologists and social psychologists approach the measurement of trust as a feature of a relationship with a specific individual, in a specific context. For instance, in a review of the most widely adopted attitudinal measures in social psychology, Wrightsman (1991: 375) highlights the “Specific Interpersonal Trust Scale” developed by Johnson-George & Swap (1982) as an exemplar since it was “designed to measure trust of another person under particular circumstances.” Likewise, sociologists invariably measure trust as one person’s expectations of positive intentions and motives of a particular person in a specific situation (Molm, Takahashi & Peterson 2000).<sup>3</sup>

If we accept that individuals can exhibit different levels of trust for different targets, then considering targets is critical to the measurement of trust. Most economic studies, however, ignore

---

<sup>3</sup> Among social psychologists, even the development of an instrument to measure an individual’s *general* propensity to trust others first involves identifying the amount of trust that individual exhibits for several targets, and only then inferring whether that individual is generally trusting of others. As Rotter (1967) notes in developing his inter-personal trust scale, “An attempt was made to sample a wide variety of social objects, so that a subject would be called upon to express his trust of parents, teachers, physicians, politicians, classmates, friends, etc.” (p. 653).

the target of trust and instead treat an individual's willingness to trust as a stable trait that can be adequately measured using a single behavioral measure with a single target.<sup>4</sup> Thus, when an experiment using the investment game reveals a laboratory subject  $i$  to exhibit a level of trust  $x_i$ , we must be cautious to note that what we have really measured is  $x_i(j)$ , where  $j$  represents the target (usually another anonymous subject in the experiment).

While attending to the target of trust is important for trust measurement in general, it is particularly important for comparing and integrating research on trust across the social sciences, especially in cases where different approaches are used to measure trust. To make a valid comparison, different measures must address the same target. For example, suppose we compare individual  $i$ 's behavior in an investment game with target  $j$  ( $s_i(j)$ ) to an attitudinal measure of trust towards target  $k$ ,  $a_i(k)$ . It is not surprising if the two measures do not correlate. To test the reliability of the two measures, we should hold the target constant, comparing  $s_i(j)$  to  $a_i(j)$ .<sup>5</sup>

The importance of identifying the target of trust is relevant to interpreting the study by Glaeser et al and provides insight into the lack of correspondence across the measures of trust in their and others' studies. The primary behavioral measures used by Glaeser et al consisted of the investment game and an "envelope drop" in which subjects stated their reservation price for a self-

---

<sup>4</sup> There are a few exceptions, though none highlights the importance of target-specificity to the extent we do here. Eckel and Wilson (2006) study how much trust differs based on the attractiveness of the target (see Scharlemann et al, 2001) and Buchan, et al (2002) demonstrate that trust is negatively related to the "social distance" between the trustor and the target. Even so, however, these studies do not note the importance of target-specificity as a general principle in the measurement of trust, nor do they explore how behavioral and attitudinal measures are related when accounting for the target of trust. Two studies make a point closer to ours. Holm and Danielson (2005) distinguish between "thick" trust (towards those with whom one has close social ties) and "thin" trust (towards people in general). Yet, by correlating "thin" measures (General Social Survey questions) with behavioral trust within potentially "thick" populations (students in the same economics classes) they do not hold constant the specific target of trust in their experiment, as we do here. Finally, Knack and Keefer (1997) note a distinction between "generalized" trust (towards people in general) and "specific" trust (towards "people one has repeated interactions with" (p. 1258)), and present an example of how target-specificity matters for relating measures of trust. However, the discussion of target-specificity is tangential in their paper.

<sup>5</sup> In fact, the social psychology, sociology, and organizational literatures further suggest that the context in which trust occurs is important. Individual  $i$  may trust  $j$  with a monetary investment (as one might with a financial advisor), but the same individual  $i$  may exhibit almost no trust towards  $j$  in other contexts, such as caring for  $i$ 's children.

addressed envelope containing money to be dropped at a specified location. Glaeser et al compared the degree of trust exhibited by subjects in these contexts to 12 different attitudinal measures, ranging from a broad General Social Survey (GSS) item to more specific questions. Of the 12 attitudinal measures, only two significantly predicted trusting behavior in either game.

However, Glaeser et al primarily tested the reliability of different measures of trust across targets. For instance, they compared the very general GSS question, which measures trust in “most people,” to behavior in an investment game in which some subjects were paired with someone they knew.<sup>6</sup> If a subject thought that people in general are trustworthy but knew that the particular matched person was not, then it is unsurprising to find no relationship between the measures.

To address such concerns, it is necessary to compare measures while holding the target constant. In fact, some comparisons between attitudinal and behavioral measures in Glaeser et al’s data support our claim that such measures are related only when the target of trust is held constant. Recall that the two behavioral trust measures in their study were an “envelope drop,” in which the targets were strangers at different locations around Boston, and the investment game, in which subjects played either against an anonymous other student – a “stranger” – or someone they knew. Of all the attitudinal measures of trust, only two correlate significantly with the behavioral measures. Both of these questions deal specifically with “strangers” as the targets of trust. Subjects who express trust in strangers in these two attitudinal questions state significantly higher reservation prices for the envelope drop (indicating greater trust). Similarly, subjects who express greater trust of strangers in the attitudinal measures send significantly more money to the second player in the investment game, *but only when playing against an anonymous opponent*. Thus, when the target of trust is held constant to “strangers” – anonymous individuals one is likely to encounter in passing – Glaeser et al’s results reveal strong positive relationships between behavioral and attitudinal data.

---

<sup>6</sup> Similarly, Karlan (2005) compared behavior in non-anonymous investment game pairings with GSS trust questions.

The experiment we describe below directly addresses the importance of the target of trust by explicitly varying the counterpart across repeated rounds of the same game. In particular, subjects in our experiment play the game with individuals at five different non-laboratory locations, each of which implicitly varies several characteristics of the target, including perceived trustworthiness. We also compare trust elicited via our behavioral measure (which we describe below) with attitudinal measures of trust, with the two types of measures collected at different points in time. We find that when accounting for variation in the target, there exists a positive relationship between attitudinal and behavioral measures.

*B. Measuring (dis)trust as the willingness to pay to avoid vulnerability*

In addition to establishing a link between behavioral and attitudinal measures by noting the importance of target-specificity, we also propose an alternative behavioral approach to measuring trust. While behavioral measures such as the investment game approach the measurement of trust by attempting to elicit the act of trusting (i.e., by observing how much money subjects send to another player), we note that the phenomenon of trust is often evident and easily measured by observing the extent to which individuals *distrust* others.

When trust is absent or incomplete, individuals frequently incur costs to mitigate their vulnerability. We view the willingness to incur such costs as an inverse indicator of trust. For example, if one individual fully trusts another, she is willing to enter into an exchange with no safeguards to mitigate her vulnerability to the other party. However, when trust is absent, individuals often pay for mechanisms – such as contracts, monitoring, or enforcement – to mitigate such vulnerability. The extent to which an individual is willing to pay for such mechanisms is inversely proportional to the level of trust exhibited towards the target. Thus, we propose to

measure trust/distrust through individuals' willingness to incur costs to mitigate vulnerability to the target.

We believe this alternative approach is valuable for a couple of reasons. First, it corresponds to how one frequently observes trust/distrust in the real world, with economic agents incurring costs to mitigate their vulnerability to others. Such costs are often identifiable and measurable. For example, managers' costly monitoring of employees, homeowners' willingness to purchase security services, and suspicious spouses' willingness to pay for surveillance are all quantifiable measures of distrust by one party towards a specific target or a specific group of individuals.

Second, our approach avoids an important problem with the investment game. In the investment game, the act of trusting involves "giving" to someone else. Thus, the key action in the most widely-used behavioral measure of trust is almost identical to the key action in the most widely-used behavioral measure of altruism or fairness, the dictator game (Hoffman et al 1994). As Cox (2004) notes, measuring how much the first mover in the investment game sends to the second mover does not "discriminate between actions motivated by trust . . . and actions motivated by other-regarding preferences characterized by altruism . . . that is not conditional on the behavior of others" (p. 262). If individuals derive utility from the mere act of giving to others (Andreoni, 1989, 1990), then the investment game is very likely to confound this motive with trust. The behavior in our measure – indicating a willingness to pay to mitigate vulnerability to another – is fundamentally distinct from the act of giving. Moreover, in the investment game, distrustful behavior (i.e., investing zero) penalizes trustworthy and untrustworthy targets alike. However, in our game the act of distrust only penalizes the other player if he intended to appropriate more than the fair share.<sup>7</sup>

---

<sup>7</sup> It is also straightforward to demonstrate that the social motives in most models of fairness or altruism (e.g., Fehr & Schmidt, 1999; Bolton & Ockenfels, 2000) apply more strongly to the trustor's action choice in the standard investment game than to the game we use here.

To formalize our procedure and present the game we use in our experiment, suppose that two individuals,  $i$  and  $j$ , can potentially enter into an exchange that yields a surplus  $W$ , that they have agreed to divide evenly.<sup>8</sup> Now, suppose that  $j$  possesses agency over the final distribution of  $W$ , meaning that he can appropriate any portion ( $x$ ) of  $W$ , leaving  $i$  with the remainder ( $W-x$ ).

The above example is one in which the extent to which  $i$  trusts  $j$  matters for whether the exchange occurs and under what conditions it occurs. If  $i$  fully trusts  $j$  to divide the surplus evenly ( $x = \frac{1}{2}W$ ), then she should agree to the exchange with no additional safeguards. However, suppose instead that  $i$  exhibits some distrust towards  $j$ , and expects him to appropriate a larger share of the surplus ( $x > \frac{1}{2}W$ ). Then  $i$  should be willing to pay some positive amount for an instrument to mitigate her vulnerability to  $j$  – or, put differently, for some mechanism that constrains  $j$ 's behavior.

In particular, suppose  $i$  can pay some amount,  $p$ , to ensure an equal division of  $W$ . Doing so eliminates  $j$ 's agency over  $W$ , and guarantees the final payments to be  $\pi_i = \frac{1}{2}W - p$  and  $\pi_j = \frac{1}{2}W$ . The maximum price that  $i$  is willing to pay to avoid vulnerability to  $j$ ,  $p^*$ , measures  $i$ 's distrust towards  $j$ . If  $i$  fully trusts  $j$  then  $p^* = 0$ , while if  $i$  fully distrusts  $j$  then  $p^* = \frac{1}{2}W$ .

Our paradigm is simple. We create the exchange above and then elicit from player  $i$  the maximum price she is willing to pay to mitigate her vulnerability to the other player ( $j$ ), or  $p^*$ . Our hypothesis is that this will accurately measure distrust, evidenced both by weak correlations with other behaviors such as altruism or risk aversion and by a strong (negative) correlation with attitudinal measures of trust when holding the target ( $j$ ) fixed.

### *C. Our Experiment*

---

<sup>8</sup> Of course, they could in principle agree to any other division, against which the trustworthiness of  $j$  can be evaluated. For example, the two agents may simply agree that  $j$  will divide the surplus “fairly” or “equitably,” taking into account inequality, costs of inputs, utility, etc. (cf. Yaari & Bar-Hillel 1984; Konow 2000; Tungodden et al 2007), as long as this allows  $i$  to form an expectation of what such a division should be.

We used the above game with an endowment ( $W$ ) of \$10. Player  $i$  was always a laboratory subject. To account for the importance of the target of trust, we varied the role of  $j$  across several non-laboratory participants and elicited reservation prices from each laboratory subject with respect to every possible target ( $p_{ij}^*$ ).

We obtained targets with varying characteristics by pairing laboratory subjects with randomly-selected field participants passing by several locations around the city of Pittsburgh.<sup>9</sup> Each laboratory subject in the role of Player 1 ( $i$ ) played the game five times, each time with a new random field participant in the role of Player 2 ( $j$ ) from one of five specified locations. Each of these games was played independently and for payment, and we counterbalanced the order of locations.

We selected five locations that varied considerably in population characteristics (socioeconomic status, ethnicity, education). To measure these characteristics and ensure that we obtained variability in the targets, field participants in the role of Player 2 completed a questionnaire eliciting several demographic variables. We also measured the perceptions of these characteristics by laboratory subjects in the role of Player 1 (discussed further below), to assess the accuracy of their expectations about the target and to explore the determinants of trust.

Finally, to compare our behavioral distrust measure to attitudinal measures and to behaviors unrelated to trust, laboratory subjects completed several other tasks. These additional tasks measured altruism towards the target, risk seeking, and trust attitudes towards the target.

### **III. Experimental Design**

---

<sup>9</sup> Other researchers simultaneously employed a similar methodology of pairing laboratory subjects with people from districts around a city (Zurich), using the investment game (Falk and Zehnder, 2007). There are important differences between our studies. For example, our research focuses on a new behavioral measure and on the relationship between behavioral and attitudinal measures when accounting for the target of trust. In fact, our primary purpose for using this design is to create the variability in targets necessary to test this relationship. In addition, we focus much more on specific characteristics of the individual target – rather than of the general population at the target location – for instance, by measuring laboratory subjects’ beliefs regarding the specific characteristics of the target individual and using those predicted characteristics to predict distrust towards the target.

We recruited laboratory subjects from the University of Pittsburgh's Experimental Economics Laboratory subject pool. Each laboratory subject participated in two separate sessions, about three to four weeks apart. To recruit field participants we went to five different locations in the city of Pittsburgh and asked people passing by, one at a time, to stop and participate in a brief experiment for pay.

*A. Laboratory subjects – Session 1*

Laboratory subjects were recruited via e-mail from a list containing students and staff at the University of Pittsburgh. Before signing up, subjects were told that they would have to attend two sessions, three to four weeks apart, and would receive the majority of their payment at the conclusion of the second session. Upon arriving, subjects received a private participant number, to be used in both sessions. Subjects were seated apart from one another and received instructions, which were handed to them and also read aloud.

**The behavioral distrust measure.** The instructions (see Appendix A) stated that laboratory subjects (Player 1) would play a game in which they would be paired with someone (Player 2) from a specific location in the city of Pittsburgh that would be specified later. The laboratory subjects saw the instructions that were to be shown to field participants.

The field participant would be told that he or she was matched with a laboratory subject in an experiment at the University of Pittsburgh, and that the pair had been allocated a combined \$10, of which \$5 was designated for each person. The field participants would then specify how much of the total wealth, including the laboratory subjects' \$5, he or she would like to keep ( $\$10 \geq x \geq \$5$ ). The field participant would receive either this amount or \$5, determined by whether the paired

laboratory subject had paid to bind the field participant's behavior (which was not known by the field participant at the time of selecting an amount to keep).

After viewing the instructions for the field participant, laboratory subjects were told that they could specify how much they were willing to pay out of their own \$5 share to not play the game and instead guarantee the remainder of their \$5. That is, they could pay an amount between \$0 and \$5 to eliminate the field participant's discretion over the division of the endowment.

Laboratory subjects were shown a "Decision Sheet" (see Appendix A) containing 51 rows, each corresponding to an amount  $0.00 \leq p \leq 5.00$ , in ten cent increments. For each row, laboratory subjects specified whether they would prefer to pay the corresponding amount or "play the game" by selecting one of two boxes. One of the rows would be selected at random and the subject's choice on that row would determine whether the field participant would have the opportunity to keep more than \$5. More precisely, if a row corresponding to an amount  $p$  was selected, then if Player 1 selected to pay that amount the payoffs would be  $\pi_1 = 5 - p$  and  $\pi_2 = 5$ , while if for that amount Player 1 selected to play the game then the payoffs would be determined by Player 2's choice of  $x$  ( $\pi_1 = 10 - x$  and  $\pi_2 = x$ ). Laboratory subjects were told that, regardless of the outcome, they would receive their earnings at the second session.

The above procedure is structurally identical to the Becker, DeGroot and Marshak (1964) lottery-elicitation procedure for incentive-compatible valuation. A subject ( $i$ ) deciding whether to trust a target ( $j$ ) should identify the maximum price he or she is willing to pay to not play the game,  $p_{ij}^*$ , by selecting to pay at any value  $p \leq p_{ij}^*$  and selecting to play the game for any value  $p > p_{ij}^*$ . Thus, the row on the Decision Sheet at which a laboratory subject switches from one column to the other indicates the most she is willing to pay to avoid vulnerability to that target.<sup>10</sup>

---

<sup>10</sup> To minimize inconsistent responses – for instance, in which a subject prefers to play the game at some amount  $p'$  but not at an amount  $p'' > p'$  – we explained to subjects that one way to complete the Decision Sheet was to draw a vertical

After receiving instructions, laboratory subjects completed a practice Decision Sheet, for which they were hypothetically paired with a field participant from the adjacent hallway outside the laboratory. An experimenter publicly drew a number from 1 to 51 to select a row (price) and explained publicly what would happen for each of the two possible choices for that row based on the behavior of the hypothetical passerby. Laboratory subjects then completed a quiz to ensure understanding of the game and procedures.

We then proceeded to the first round of the game. Each laboratory subject received a Decision Sheet corresponding to one of the five locations in the city of Pittsburgh. The top of the Decision Sheet included a written description of the location (the name of the neighborhood and an intersection for two major roads) and the back included both a map of Pittsburgh with an arrow pointing to the location and a picture of the location. Since laboratory subjects were to receive the five locations in different orders, the first location differed privately. Laboratory subjects were not told in advance how many times they would play the game.

Each laboratory subject selected a response for each of the 51 rows on the Decision Sheet, and an experimenter collected the Decision Sheet. Laboratory subjects then were told they would play the game again, with a new location indicated on the new Decision Sheet. This procedure continued until all subjects had played five times, once for each location.

**Locations.** The five locations consisted of two universities and three residential/commercial neighborhoods. The university locations were the student centers at the University of Pittsburgh and Carnegie Mellon. The neighborhoods consisted of one upper-income neighborhood (Shadyside), one middle-income neighborhood (the Southside), and one lower-income neighborhood (the Hill District). The three specific non-university locations were commercial areas, but all three were

---

line through several boxes down to the maximum amount they would pay to not play the game and to then draw another line through the other column for the remaining rows. However, subjects were not required to do so.

within a block of residential areas. Appendix Table B2 provides summaries of 2000 U.S. Census demographic data for each non-university location.

**Risk seeking.** After playing the game five times, laboratory subjects completed a behavioral task measuring risk attitudes. We used the same task as Schechter (2007) because she found a relationship between risk seeking and trust measured by the investment game.

Laboratory subjects received \$4 to start. They decided how much of that \$4 to bet on a die roll. Their earnings from the task were any amount they did not bet, plus any earnings from the die roll. Their earnings from the die roll were the amount bet times a multiple determined by the die roll. For each possible die roll (1, 2, 3, 4, 5, or 6) the corresponding multiple was 0, 0.5, 1, 1.5, 2, or 2.5. Subjects received payment for this task at the end of the first session.

**Altruism.** To measure possible altruism or interpersonal concern for each target, laboratory subjects played a dictator game with a new field participant from each of the five locations. For each of the five locations, laboratory subjects allocated \$10 between themselves and a new field participant from that location. Subjects were informed that one of these locations would be selected at random at the end of the session. For the selected location, the experimenters would recruit a new person passing by with whom to match the subject (someone who had not played the distrust game) and would give this person the allocated portion of the \$10. Any portion of the \$10 not allocated to this second field participant would be paid to the laboratory subject at the second session.

**Perception of targets' demographic characteristics.** For the final task in the first session, laboratory subjects received five copies of a one-page demographic questionnaire that was to be completed by the same field participant at each location with whom they would be paired for the distrust game. Laboratory subjects were told to complete each questionnaire, attempting to match the responses of the paired field participant at the specified location.

At the end of the session, one location and one question would be selected at random. If the laboratory subject matched the response on this item to that of the paired field participant, the laboratory subject would receive an additional \$5 in the second session.

**End of first session.** Once laboratory subjects completed all the above tasks, an experimenter conducted several random draws to determine outcomes for the various measures. Following public announcement of the random draws, laboratory subjects were paid a \$6 participation fee and their earnings from the risk seeking task.

### *B. Laboratory subjects – Session 2*

After approximately three weeks, we contacted via e-mail every laboratory subject that participated in Session 1 and provided them with a list of several dates and times from which they could select a second session.<sup>11</sup> Upon arriving at one of these sessions, laboratory subjects received a set of questionnaires and were asked to proceed sequentially through each sheet.

**Attitudinal trust measures.** The first part of the questionnaires asked laboratory subjects to state agreement or disagreement with several statements reflecting trust attitudes towards people at the five locations from which we selected field participants. Each sheet corresponded to a particular location, which was clearly indicated at the top of the sheet. For each location, we asked three sets of questions, which varied in how specific they were towards the precise target and the context of the distrust game.

The first set of questions dealt with *trust towards the specific target and context of the game* from Session 1. Laboratory subjects rated agreement on a four-point scale (strongly disagree, disagree somewhat, agree somewhat, strongly agree) with five statements about the field participant

---

<sup>11</sup> Every laboratory subject except one returned for one of the second sessions. The one laboratory subject who was unable to attend any of the scheduled sessions completed the second part individually at one of our offices.

with whom they were paired at the specified location.<sup>12</sup> The statements are representative of the kind frequently used by social psychologists, sociologists, and organizational researchers to measure trust (Cummings & Bromiley 1996; Mayer & Davis 1999).

The next set of questions dealt with *trust towards related targets and contexts*. First, laboratory subjects were asked to imagine that they had left their wallet containing \$60 in a store at the specified location. They then rated, on the same four-point scale, both how likely it would be that someone at that particular location would turn their wallet in to the store and, conditional on the wallet being retrieved, how likely it would be that the money would be in the wallet. Next, laboratory subjects were asked to imagine that they were considering purchasing tickets to a sold-out event advertised for sale in the local newspaper, and that doing so required mailing a check to a post office box at the specified location. Subjects again indicated two responses on the four point scale: how likely they would be to mail the check and, conditional on mailing the check, how likely the other person would be to mail the tickets.

Finally, we asked one question regarding the *general trustworthiness* of individuals at the particular location. The question read, “I believe that people at [specified location] are generally trustworthy” and subjects again rated agreement using a four-point scale.

The questions all generally apply to people at the indicated location. Therefore, we expect the behavioral measure of distrust towards people at that location to be negatively related to the three attitudinal trust measures. Also, since we vary the degree to which the attitudinal measures are target and context specific, we predict that the correlation will be strongest for the attitudinal measure that deals with trust in the *specific target/context* and weakest for the attitudinal measure that deals with *general trustworthiness*.

---

<sup>12</sup> Questions asked whether the target was someone likely to “be trustworthy,” “do the right thing,” “treat me fairly,” “take into consideration my interests,” and “take advantage of my vulnerability.” We reverse-coded the last item.

**Familiarity with location.** We then elicited laboratory subjects' familiarity with each location. Subjects indicated how familiar they were with the location, how often they had visited the location in the past year, and the number of people from the location with whom they regularly interacted.

**Demographics and general background.** In the final part, subjects provided their own demographic information (identical to the items asked of field participants). We also asked several questions related to subjects' general background and specific experiences. These questions were identical to those used by Glaeser et al (2000) and found to be significantly related to the behavioral measures of trust in their study. This included one measure of general trust in strangers ("You can't count on strangers anymore").

**Feedback and final payment.** After completing the questionnaires, each laboratory subject received a sheet with information on the outcomes for all five instances of the distrust game played with field participants and whether that subject obtained the bonus for correctly predicting a demographic response of the field participant. We then paid laboratory subjects their earnings from the five games, the dictator game, and any bonus for correctly predicting demographic characteristics. Laboratory subjects earned \$43.51 on average across both sessions.

### *C. Field participants*

Before the second laboratory session, the experimenters went to each of the five locations in teams of between 1 and 3 experimenters. We visited each location in the afternoon on a weekday and more than once. Upon encountering a passerby, the experimenter identified himself as a university researcher and asked if the individual would participate in a brief experiment, lasting about two minutes, in return for at least \$5 and possibly more.

**Behavioral untrustworthiness measure.** If a person agreed to participate, the experimenter then read a brief instruction sheet, which was identical to the one that had been shown to laboratory subjects for illustration purposes. We told field participants that they were paired with a laboratory subject at the University of Pittsburgh, that the pair had been allocated \$10, and that the field participant would receive his or her \$5 and could also claim up to \$5 of the laboratory subject's allocation. Field participants were told that they would receive either this claimed amount or \$5, based on factors outside of their control. Field participants then indicated the amount they wished to keep ( $\$5 \leq x \leq \$10$ ) on a decision sheet. They then found out whether they would receive either this claimed amount or \$5, with the outcome dependent on whether or not the paired laboratory participant had opted to play the game (determined by the laboratory participant's choice on the selected row during Session 1).

**Demographic questionnaire.** While the experimenter prepared the payment, field participants completed a one-page questionnaire with five multiple-choice demographic questions (age, gender, ethnicity, education, and socio-economic status). These were the same questions that laboratory participants had tried to predict in their first session.

**Dictator allocations.** Depending on the outcome of the random draws in prior laboratory sessions, we also distributed dictator game earnings at some locations. This occurred only after collecting distrust game and demographic data at that location. To distribute dictator game earnings, we approached passersby, identified ourselves as university researchers, and asked if they would accept an envelope that may or may not contain money, based on the outcome of a previous experiment. If someone agreed, then that person drew an envelope containing a small sheet that briefly described the dictator game and the corresponding amount shared (if any).

## IV. Results

We obtained data from 60 laboratory subjects and 300 field participants (60 at each location).<sup>13</sup> Table B1, in Appendix B, provides self-reported demographic information on laboratory and field participants at each location, as well as laboratory subjects' predictions of field participants' responses to the same demographic questions.

Our primary dependent variable, distrust, consists of the price at which a laboratory subject switched from preferring to pay to avoid the game to playing the game. In the majority of cases, laboratory subjects switched from the column “pay, don't play” to the column “don't pay, play” only once. In these cases the amount corresponding to the row at which a laboratory subject switched to the “don't pay, play” column was that subject's threshold price for that location ( $p_{ij}^*$ ).

In a few cases, laboratory subjects indicated inconsistent responses by changing between columns more than once on a Decision Sheet (for instance, by indicating a willingness to pay \$2.00 and \$0.50 but not \$1.00). Out of 300 completed Decision Sheets (by 60 laboratory subjects), this occurred a total of 10 times (by 4 different laboratory subjects). Responses of this sort do not provide us with a clear threshold price. Nevertheless, we can still use information on the Decision Sheet to infer a roughly equivalent statistic measuring how many times, and for what values, the laboratory subject selected boxes in each column.<sup>14</sup> This statistic tells us something about how

---

<sup>13</sup> We also distributed 60 dictator recipient envelopes, but collected no data from these field recipients.

<sup>14</sup> More precisely, we represent subject  $i$ 's choice to either pay (1) or not pay (0) in each of the 51 rows of the Decision Sheet by  $x_{ik} \in \{0,1\}$ , where  $k = \$0.00, \$0.10, \dots, \$5.00$ . Then, we compute the following statistic as an imputed measure of the cutoff, based on the number of choices in each column:

$$\hat{p}_{ij}^* = \$0.10 + 2 \left( \frac{\sum_k x_{ik} k}{\sum_k x_{ik}} \right)$$

For a subject choosing “consistently” (only switching between columns once), this formula returns the first value in the “don't pay, play” column.

We also conducted our analysis excluding the individuals who provided at least one inconsistent set of responses. The results are generally unchanged by omitting these four individuals. We discuss specifically how they change in instances where this is the case.

much that laboratory subject is willing to pay to avoid vulnerability to the field participant and is comparable to the threshold price used for the rest of the analysis.

Figure 1 shows the distribution of these maximum buy out prices. They range from \$0 (complete trust, five instances) to \$5 (complete distrust, two instances), with a mean of \$2.61, median of \$2.60, and standard deviation of \$0.97. While there is a central tendency in the data, over half of the responses (51 percent) lie outside of the interval [\$2, \$3].

#### *A. Individual-level differences among laboratory subjects*

We first explore the distribution of average maximum buy out prices by individuals. For each laboratory subject we compute the average maximum price across all five locations ( $\hat{p}_i^* = (\sum_{j=1}^5 p_{ij}^*)/5$ ). We can interpret this as an individual's propensity to distrust passersby at public locations in the city of Pittsburgh. This propensity ranges from highly trusting (\$0.51) to highly distrusting (\$4.68), with a mean of \$2.61 and standard deviation of \$0.84. Figure 2 presents the distribution of individual-level average prices. Over 40 percent (25 of 60) of these individual-level average prices lie outside the interval [\$2, \$3].

Table 1 explores what factors determine the individual-level propensity to distrust. The first regression explores the impact of laboratory subjects' own demographic characteristics. *Age* is coded as 0 if a laboratory subject responded in the "18-24" category and 1 otherwise. *Female* is included as a binary variable. *Race* is included as two binary variables, one for "Black" and another for all categories other than "Black" or "White." *Education* is also a binary variable, coded as 0 if a laboratory subject responded "some college" (no laboratory subjects reported lower levels) and 1 otherwise (indicating attainment of either a bachelor's or graduate degree). Social / economic status (SES) is included as two binary variables, one for "Lower"/"Lower-Middle" and another for

“Upper”/“Upper-Middle.” Finally, *Only child* is a binary variable indicating whether a laboratory subject reported having any siblings.

The remaining regressions in Table 1 test several additional variables measuring past experiences and behaviors, many of which were also used by Glaeser et al (2000). *Door unlocked* is a binary variable indicating whether a laboratory subject reported ever leaving his or her door unlocked.<sup>15</sup> *Lend money* and *Lend possessions* are laboratory subjects’ responses on a 1 to 4 scale indicating the frequency with which they engage in these behaviors with friends, with higher numbers corresponding to greater frequency.<sup>16</sup> Following Glaeser et al, we construct a *Trusting behaviors index* by standardizing and summing the previous three variables. Finally, *Benefit from stranger* indicates a positive response to the question, “Have you ever spontaneously benefited from the generosity of someone you never knew before?”

Five variables predict distrust reliably at the individual level. Older laboratory subjects, females, and non-white/non-black subjects are generally less distrusting; more educated laboratory subjects are more distrusting.<sup>17</sup> Laboratory subjects who report having previously benefited from a stranger’s generosity are less distrusting. This is consistent with Glaeser et al, who interpret this as reflecting individuals’ expectations of the likely behavior of a stranger. Interestingly, we find almost

---

<sup>15</sup> The possible responses to the question, “How often do you leave your door unlocked?” were “Very often” (3%), “Often” (2%), “Sometimes” (10%), “Rarely” (42%), and “Never” (43%). The final response category was coded as 0 and all others as 1.

<sup>16</sup> The questions ask, “How often do you lend [money / personal possessions] to friends?” Possible responses were, “More than once a week” (=3), “About once a week” (=2), “About once a month” (=1), and “Once a year or less” (=0). Mean responses were 0.85 (money) and 1.25 (possessions). Glaeser et al (2000) used a slightly different coding, in which higher numbers corresponded to less frequency. For easier interpretation, we code responses so that higher numbers correspond to greater frequency.

<sup>17</sup> If we conduct the analysis excluding the four subjects who responded inconsistently at least once (see footnote 14), the coefficient on *Age* becomes larger in both significance and magnitude.

no relationship between distrust and self-reported trusting behaviors, which is inconsistent with Glaeser et al's results.<sup>18</sup>

Table 2 explores the relationship between individual-level distrust and three other individual-level behaviors or attitudes. First, we construct a broad measure of general individual-level trust attitudes, *Overall trust attitude*, using all of a subject's responses to the attitudinal trust questions.<sup>19</sup> *Altruism* is the average amount that a laboratory subject shared in the five dictator games. *Risk attitude* is the amount bet by a laboratory subject on the die roll. For these regressions, we standardize all three measures by subtracting the mean and dividing by the standard deviation.

We find little relationship between risk and distrust. Altruism and overall trust attitude have similar relationships with distrust. The coefficients indicate less behavioral distrust when a subject exhibits greater attitudinal trust or greater altruism, by a roughly \$0.20 lower maximum buy-out price for each standard deviation. Both coefficients are significant at  $p = 0.06$ .<sup>20</sup> *Overall trust attitude* and *Altruism* correlate quite highly ( $0.44, p < 0.001$ ), indicating a relationship between attitudinal measures of trust and behavioral measures of altruism.

### *B. Determinants of target-specific trust*

Table 3 presents summary statistics, by location, of behaviors and attitudes from both laboratory subjects and field participants. The variables in the table are the average buy-out price

---

<sup>18</sup> Lending money to friends is significant in a few regressions, but the coefficient is relatively small. Moreover, the sign of the coefficient is contrary to what we would expect: laboratory subjects tend to be more distrusting the more frequently they lend money to friends.

<sup>19</sup> Recall that for each of the five locations, subjects provided 10 attitudinal responses concerning their beliefs and expectations regarding trustworthiness of people at that location. We first average an individual's responses, across the five locations, separately for each of these questions. We then add together all of the 10 averaged responses and also the subject's agreement with the statement, "You can't count on strangers anymore," (we reversed the sign of the three items for which agreement with a statement reflects distrust). Later in the paper, we more thoroughly explore the relationship between specific attitudinal questions and trust for specific targets.

<sup>20</sup> Excluding the four individuals who provided inconsistent responses increases the significance and magnitude of the coefficient for Age, slightly increases the magnitude and the significance of the coefficient for Overall trust attitude ( $p=0.04$ ), and slightly decreases the magnitude and significance of the coefficient for Altruism ( $p=0.12$ ).

(our measure of distrust by the laboratory subject towards the field participant), amount shared in the dictator game (altruism by the laboratory subject towards a different field participant), the three attitudinal measures of trust by the laboratory subject towards people at that location, and average amount kept in excess of \$5 (untrustworthiness) by the field participant. The bottom of the table reports non-parametric (Wilcoxon) tests of pair-wise differences between locations.

Average buy-out prices (distrust) generally exhibit the same pattern as average amount kept (untrustworthiness). Hill District field participants were generally both the most distrusted and the most untrustworthy, while those in Shadyside were generally the least distrusted and the least untrustworthy. However, average buy-out prices vary less by location than excess amounts kept, and comparing the distributions of buy-out prices and excess amounts kept separately for each location (using a rank-sum test), we reject equality of distributions for every location.

On average, laboratory subjects generally distrusted too much, as mean buy-out prices are usually higher than excess amounts kept. In terms of expected value, laboratory subjects paid \$0.60 per location (\$3.00 in total across locations) too much to distrust (\$2.61 average buy-out price across locations vs. \$2.01 average excess amount kept). For the first four locations in Table 3, they paid on average \$0.89 too much, while for the Hill District they paid \$0.58 too little.<sup>21</sup>

Comparing across locations (as opposed to across individuals as in Table 2), the average amount shared in the dictator game (altruism) is generally *positively* related to distrust. That is, the populations generally regarded as untrustworthy (high buy-out prices) are also the ones with which the most is shared. This contrasts with the investment game, where high trust generally corresponds to high altruism (e.g., Cox 2004; Ashraf et al 2006).

---

<sup>21</sup> If we use the average excess amount kept at a location as a standard against which to compare whether particular laboratory subjects distrusted too much or too little, individual subjects generally distrusted too much (222 of 300 buy-out prices are higher than the corresponding average excess amount kept for that location). The same pattern holds separately for all locations (between 49 and 55 out of 60) except for the Hill District, where subjects distrusted too little (only 16 of 60 buy-out prices are higher than \$3.38).

The attitudinal measures exhibit similar patterns to buy-out prices. For example, the least trusted group for all attitudinal measures is the Hill District, which is also the group most distrusted and most untrustworthy in the context of the game. Shadyside is generally trusted attitudinally, mirroring low distrust and low untrustworthiness in the game. However, there are also differences. Southside people are generally not trusted in the attitudinal measures but are distrusted relatively little in the game. They also prove to be relatively trustworthy in the game.

There is also variation between the different attitudinal measures. For instance, people at the University of Pittsburgh are not considered particularly trustworthy in the specific context of the game (the average rating of 0.13 is significantly lower than the ratings for both CMU and Shadyside). However, in related contexts they are considered much more trustworthy.

To more closely examine the relationship between these behaviors and attitudes, we estimate regressions in which a laboratory subject's maximum buy-out price for a location (in contrast to the average buy-out price across locations used in Tables 2 and 3) is the dependent variable. These regressions provide target-specific analyses of the determinants of trust. Since the models include repeated observations for each laboratory subject, we include subject fixed effects. These regressions are reported in Tables 4 and 5.<sup>22</sup>

The first regression in Table 4 explores the relationship between predicted demographic characteristics of the field participant, by the laboratory subject, and the laboratory subject's distrust towards the field participant. Laboratory subjects' predictions of field participants' demographic characteristics were generally accurate (see Table B1 in Appendix B).<sup>23</sup> As explanatory variables,

---

<sup>22</sup> One participant failed to answer the predicted demographics for one location (University of Pittsburgh). Therefore, we have 299 predicted demographics questionnaires completed by 60 laboratory participants. Another participant did not complete one of the attitudinal questionnaires. This means that we have complete data (behavioral and attitudinal measures) for 298 targets, by 60 laboratory participants.

<sup>23</sup> In most cases where there is a clear modal response by field participants at a location (at least 20% higher than the next most frequent response), the modal predicted response is the same. The three exceptions are easy to interpret. First, laboratory subjects mispredicted the modal ethnicity response at CMU (68% predicted Asian while the modal response

we construct predicted demographic variables in the same way as we did for laboratory subjects' own demographic characteristics (see Table 1), but also include one additional variable to account for greater heterogeneity among field participants. *Predicted Education* now comprises two binary variables – one indicating a predicted response of “some or no high school” or “high school degree or equivalent” (no college) as well as the binary variable we used previously indicating the attainment of at least a bachelor’s degree (as in Table 1, the omitted category is “some college”).

The second regression includes several binary variables indicating whether a laboratory subject predicted the target to have the same demographic characteristics as him or herself. Each of these binary variables is equal to 1 if the predicted response exactly matches the laboratory subject’s own response and 0 otherwise.

As the first three regressions reveal, laboratory subjects exhibit less distrust towards older targets and more distrust towards targets with lower socio-economic status. Moreover, they are less distrusting of others of their own age and more distrusting of others of their own socio-economic status, though the effect of same age is modest in both magnitude and significance.<sup>24</sup>

The fourth and fifth regressions explore the effect of familiarity with the location. For each location, laboratory subjects rated how familiar they were with the location, how many people they knew at the location, and how often they frequented the location. We summed these variables and

---

was White (55%)). This likely reflects stereotypes of CMU’s population consistent with representativeness bias (Tversky and Kahneman, 1974). The two other exceptions occurred when laboratory subjects predicted modal socio-economic status responses other than “Middle,” which was the modal response for every location. These differences are perhaps similarly due to stereotyping, but are also possibly due to reporting bias on the part of field participants (see the discussion in Results subsection C). Note, however, that in each case the distributions of predicted socio-economic status responses move in the same direction, though excessively, as the actual responses.

<sup>24</sup> We also explored overall predicted similarity by summing the five binary variables – thus providing a measure of how many shared characteristics the laboratory subject expected to have with the target. If we replace the five binary variables with this single variable, the coefficient is both small in magnitude and statistically insignificant.

then standardized the resulting variable to construct a *Familiarity Index*. There is no significant relationship between this variable and buy-out prices.<sup>25</sup>

The next two regressions include the amount that a laboratory subject shared in the dictator game with a different person at that location. The amount shared is standardized. There is a very small, statistically insignificant, positive relationship between distrust towards a field participant and the amount shared with a similar person from the same location. Therefore, while we earlier found a relationship between subjects' overall propensities to share (altruism) and distrust when pooling across locations, we find virtually no relationship between altruism and distrust towards specific targets. This finding is a particularly vivid illustration of how analysis at the level of the target of trust can materially alter the findings.

Table 5 explores the relationship between the three attitudinal measures of trust towards targets at that specific location and target-specific distrust. Recall that the first measure (*specific target/context*, based on five items) dealt with trust towards the specific target in the context of the game, the second measure (*related target/context*, four items) dealt with trust in similar hypothetical contexts towards related targets, and the final measure (*general trustworthiness*, one item) dealt with overall perceptions of the trustworthiness of people at that location. All three measures are significantly (negatively) related to behavioral distrust in the game. As we predicted, the relationship is strongest in both magnitude and significance for *specific*, less so for *related*, and least so for *general*. While all three measures continue to be negatively related to distrust (with comparable magnitudes) when we include control variables, the relationship for *General* is no longer statistically significant ( $p = 0.11$ ).

---

<sup>25</sup> If we run separate regressions for each of the three familiarity variables, none of the coefficients is statistically significant.

Table 6 explores the relationship between the distrust for targets at specific locations and the average excess amount kept at those locations (i.e., average location untrustworthiness). In the first regression, there is a significant relationship between distrust and actual untrustworthiness. This relationship persists, more weakly, after we introduce the predicted demographic and attitudinal variables that we found to be significant predictors of distrust in Table 4. Thus, there is some element of distrust towards the field participants (as measured by maximum buy out prices in the game) that is related to their actual (average) untrustworthiness and is distinct from what can be explained by these predicted demographic characteristics alone.

The remaining regressions in Table 6 add the three attitudinal variables. The first two attitudinal variables predict behavioral distrust, beyond what is explained by the predicted demographic variables and actual average untrustworthiness. For the third (general) attitudinal variable, the relationship is weaker ( $p=0.12$ ). These results indicate that there is a significant component of our behavioral distrust measure that is related to attitudinal trust measures, even after controlling for important perceived and actual target characteristics.

### *C. Determinants of untrustworthiness*

Table 7 explores how self-reported demographic characteristics of the field participants relate to untrustworthiness. The dependent variable is how much money, in excess of \$5, a field participant kept. The explanatory variables are coded in the same way as in Tables 4 through 6, though here we use the actual self-reported responses of the field participants rather than the laboratory subjects' predictions of targets' demographic characteristics.<sup>26</sup>

---

<sup>26</sup> Since there exists potential censoring on both the left- and right-hand side of the dependent variable, at \$0 and \$5 respectively, we also conducted identical tobit regressions. The results do not change between the two models.

Older people are generally less untrustworthy (more trustworthy),<sup>27</sup> as are females. People with less self-reported education (no college) are generally more untrustworthy as are people with higher socio-economic status. When we do not include binary variables for each location, the coefficient for *Black* is large and statistically significant, but when we include binary location variables (omitting Pitt), the magnitude and statistical significance are both sharply reduced. Thus, what appears to be a race effect on untrustworthiness in the first regression may more accurately reflect a location effect. Finally, some of the significant relationships in Table 7 are consistent with laboratory subjects' expectations. The coefficients for *Age* and *Gender* have the same sign as those in Table 4.

Surprisingly, high socio-economic status is related to untrustworthiness.<sup>28</sup> One explanation for this puzzling relationship and for the apparent inconsistency of the *Black* variable in Table 7 lies in the possibility that SES is subject to high levels of self-report bias, particularly for the Hill District (see footnote 23). For example, as Appendix B Table B1 reveals, the modal self-reported SES for the Hill District was "Middle" and 20 percent of responses indicated "Upper-middle" or "Upper." Thus, over half of Hill District respondents self-report SES of at least middle. This is surprising, given that the Hill District is economically depressed. For example, 2000 U.S. Census data (summarized in Table B2 and Figure B2 in Appendix B) indicate that median family income in the Hill District is between \$20,000 and \$24,999, which is well below median family income in the Southside (\$30,000 to \$34,999) or Shadyside (\$75,000 to \$99,999). Therefore, field participants'

---

<sup>27</sup> If we use two binary variables to measure age (25-44 and above 44, with below 24 again as the omitted category), the effect of age is largely linear – both coefficients are negative, but the first is roughly half as large in both significance and magnitude as the second.

<sup>28</sup> To explore whether this relationship is driven largely by one location, we calculated the average amount kept by the target, by self-reported SES, separately for each location. This trend is generally true across locations – those reporting upper or upper-middle class status keep the most money for themselves at the University of Pittsburgh, CMU, Southside, and the Hill District.

SES responses might tend to be biased upwards – especially in low-SES communities like the Hill District – or might indicate status relative to others within their own community.<sup>29</sup>

To explore the possibility that response bias in the Hill District drives the significance of the Black and SES (Upper) variables, we re-ran the first regression, excluding the Hill District (model 3). When we restrict our analysis to locations other than the Hill District, the coefficients on Black and SES (Upper/Upper-middle) are no longer statistically significant.<sup>30</sup> Thus, it is likely that the earlier significance of Black and SES (Upper/Upper-middle) are driven largely by circumstances involving the Hill District, such as significant response bias in self-reported SES.<sup>31</sup>

## **V. Conclusion**

Although trust has long been a concept of interest in the social sciences, quantifying and measuring such an elusive phenomenon has proven challenging. Different methods and approaches are used across disciplines. Recently, researchers have begun to compare behavioral approaches that dominate economics with attitudinal approaches favored in other social sciences. We contribute to this effort by showing a relationship between attitudinal and behavioral measures, and by proposing a new behavioral paradigm for measuring trust that is based on measuring distrust, or individuals' willingness to incur costs to mitigate vulnerability to others.

Our findings contrast sharply with comparisons of behavioral and attitudinal measures reported in Glaeser et al. (2000) and other similar studies. Whereas Glaeser et al. observed little correspondence between behavioral and attitudinal measures, we find that the two approaches to

---

<sup>29</sup> Research in sociology has similarly noted the difficulties involved with obtaining unbiased reports of SES (see for instance Haug & Sussman 1971; Evans et al. 1992)

<sup>30</sup> While most people identifying their ethnicity as “Black” are from the Hill District (57), we still have a significant proportion of “Black” responses by field participants from other locations (32).

<sup>31</sup> Recent work by Pope and Snyder (2008) finds that people discriminate against Blacks in on-line lending and that repayment for this group is indeed lower. They conclude, similarly to us, that socio-economic status might play an important role both in distrust towards this minority group and in their trustworthiness in these contexts.

measuring trust are related when holding constant the target of trust. We also note that Glaeser et al's data contains relationships between attitudinal and behavioral trust measures, when accounting for the target.

Our results also contrast with previous research that utilizes the investment game as a behavioral measure of trust. Whereas the investment game has been shown to be confounded with risk seeking (Schechter 2007) and preferences for altruism or fairness (Cox 2004; Ashraf et al 2006), the behavioral measure reported here shows no relationship with risk, and very little relationship with altruism (i.e., a modest relationship as a subject-specific trait and none across specific targets). Moreover, our new measure of trust demonstrates predictive validity – laboratory subjects accurately discriminate trustworthiness based on perceived characteristics of targets.

We believe this study makes two important advances to the measurement of trust in economic research. First, by explicitly accounting for the target we incorporate an important factor into the measurement of trust. Our results are consistent with the view, widely held in other social sciences, that trust is often best conceived as a property of specific trustor-trustee interactions, as opposed to a stable, relatively invariant trait of the person placing trust. Second, our focus on the cost of distrust limits the potential for confounding trust with other closely related concepts. The emphasis on paying to avoid being vulnerable to another has the further advantage of external validity, as distrust is often evidenced by the extent to which individuals are willing to incur the cost of safeguards to mitigate vulnerability.

Of course, our findings are also preliminary and await further work before more definitive conclusions can be drawn. For example, as Glaeser et al demonstrate, there are clearly cases in which attitudinal and behavioral measures of trust are unrelated. Here, we motivate and present a case in which the two types of trust measures should be related. Thus, understanding when different

approaches to measuring trust are related, and how this interacts with target- and context-specificity, is an important research agenda for developing trust as a useful predictive tool for economics and social science more broadly.

Our behavioral distrust measure is useful in that it avoids problems associated with the investment game and corresponds to how distrust is frequently observed in actual economic exchanges. The principle underlying our laboratory procedure could also be used to measure distrust in other, non-laboratory settings – for example by measuring how much individuals spend to safeguard themselves against others’ moral hazard in economic settings. However, it should also be tested rigorously – as has been the case with the investment game – to attempt to identify instances in which it does not measure distrust, or where it measures unrelated concepts.

## References

- Andreoni, James. 1989. “Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence.” *Journal of Political Economy*, 97(6): 1447-1458.
- Andreoni, James. 1990. “Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving.” *Economic Journal*, 100: 464-477.
- Arrow, Kenneth J. (1974), *The Limits of Organization*. Norton: New York.
- Ashraf, Nava, Iris Bohnet, and Nikita Piankov. 2006. “Decomposing Trust and Trustworthiness.” *Experimental Economics*, 9:193-208.
- Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak. 1964. “Measuring Utility by a Single-Response Sequential Method.” *Behavioral Science*, 9(3): 226-232.
- Berg, Joyce, John W. Dickhaut, and Kevin A. McCabe. 1995. “Trust, Reciprocity, and Social History.” *Games and Economic Behavior*, 10(1): 122-142.
- Blau, Peter M. 1964. *Exchange and Power in Social Life*. New York: Wiley.
- Bolton, Gary E., and Axel Ockenfels. 2000. “A Theory of Equity, Reciprocity, and Competition.” *American Economic Review*, 90(1): 166-193.

- Buchan, Nancy R., Rachel T. Croson, and Robyn M. Dawes. 2002. "Swift Neighbors and Persistent Strangers: A Cross-Cultural Investigation of Trust and Reciprocity in Social Exchange." *American Journal of Sociology*, 108(1): 168-206.
- Capra, C. Monica, Kelli Lanier and Shireen Meer. 2007. "Attitudinal and Behavioral Measures of Trust: A New Comparison." Working paper.
- Cook, John, and Toby Wall. 1980. "New Work Attitude Measures of Trust, Organizational Commitment and Personal Need Nonfulfillment." *Journal of Occupational Psychology*, 53(1): 39-52.
- Cox, James C. 2004. "How to Identify Trust and Reciprocity." *Games and Economic Behavior*, 46: 260-281.
- Cummings, L. L., and Philip Bromiley. 1996. "The Organizational Trust Inventory (OTI): Development and Validation." In *Trust in Organizations: Frontiers of Theory and Research* ed. Roderick M. Kramer and Tom R. Tyler, 261-287. Thousand Oaks, CA: Sage.
- Currall, Steven C., and Timothy A. Judge. 1995. "Measuring Trust between Organizational Boundary Role Persons." *Organizational Behavior and Human Decision Processes*, 64(2): 151-170.
- Danielson, Anders J. and Hakan J. Holm. 2007. "Do You Trust Your Brethren? Eliciting Trust Attitudes and Trust Behavior in a Tanzanian Congregation." *Journal of Economic Behavior and Organization*, 62: 255-271.
- Eckel, Catherine C. and Rick K. Wilson. 2004. "Is Trust a Risky Decision?" *Journal of Economic Behavior and Organization*, 55:447-465.
- Eckel, Catherine C. and Rick K. Wilson. 2006. "Judging a Book by its Cover: Beauty and Expectations in a Trust Game." *Political Research Quarterly*, 59(2): 189-202.
- Evans, M.D.R., Jonathan Kellye, and Tamas Kolosi. 1992. "Images of Class: Public Perceptions in Hungary and Australia," *American Sociological Review*, 57(4): 461-482.
- Falk, Armin and Christian Zehnder. 2007. "Discrimination and In-group Favoritism in a Citywide Trust Experiment." University of Zurich, IEW-Working Paper No. 318.
- Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics*, 114(3): 817-868.
- Glaeser, Edward L., David I. Laibson, Jose A. Scheinkman, and Christine L. Soutter. 2000. "Measuring Trust." *Quarterly Journal of Economics*, 115(3): 811-846.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2004. "Does Local Financial Development Matter?" *Quarterly Journal of Economics*, 119(3): 929-969.

Haug, and Sussman. 1971. "The Indiscriminate State of Social Class Measurement," *Social Forces*, 49(4): 549-563.

Hoffman, Elizabeth, Kevin McCabe, Jason Shachat and Vernon Smith. 1994. "Preferences, property rights and anonymity in bargaining games." *Games and Economic Behavior*, 7: 346-380.

Holm, Hakan J. and Anders Danielson. 2005. "Tropic Trust versus Nordic Trust: Experimental Evidence from Tanzania and Sweden." *The Economic Journal*, 115: 505-532.

Johnson-George, C., and W. Swap. 1982. "Measurement of Specific Interpersonal Trust: Construction and Validation of a Scale to Assess Trust in a Specific Other." *Journal of Personality and Social Psychology*, 43(6): 1306-1317.

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

Knack, Stephen, and Philip Keefer. 1997. "Does Social Capital Have an Economic Payoff? A Cross-Country Investigation." *Quarterly Journal of Economics*, 112(4):1251-1288.

Konow, James. 2000. "Fair Shares: Accountability and Cognitive Dissonance in Allocation Decisions." *American Economic Review*, 90(4): 1072-1091.

La Porta, Rafael, Florencio Lopez-de-Silanes, Andrei Shleifer, and Robert W. Vishny. 1997. "Trust in Large Organizations." *American Economic Review*, 87(2): 333-338.

Lewicki, Roy J., and Barbara B. Bunker. 1996. "Developing and Maintaining Trust in Work Relationships." In *Trust in Organizations: Frontiers of Theory and Research* ed. Roderick M. Kramer and Tom R. Tyler, 114-139. Thousand Oaks, CA: Sage.

Lewis, J. David, and Andrew Weigert. 1985. "Trust as a Social Reality." *Social Forces*, 63: 967-985.

Luhmann, Niklas. 1979. *Trust and Power*. New York: John Wiley.

Mayer, Roger C., and James H. Davis. 1999. "The Effect of the Performance Appraisal System on Trust for Management: A Field Quasi-Experiment." *Journal of Applied Psychology*, 84(1): 123-136.

McEvily, Bill, Vincenzo Perronne, and Akbar Zaheer. 2003. "Trust as an Organizing Principle." *Organization Science*, 14: 91-103.

McEvily, Bill, and Marco Tortoriello. 2007. "Measuring Trust in Organizational Research: Review and Recommendations." Working Paper, University of Toronto, Rotman School of Management.

Molm, Linda D., Nobuyuki Takahashi, and Gretchen Peterson. 2000. "Risk and Trust in Social Exchange: An Experimental Test of a Classical Proposition." *American Journal of Sociology*, 105(5): 1396-1427.

Pope, Devin G. and Justin R. Snyder. 2008. "What's in a Picture? Evidence of Discrimination from Prosper.com." Working Paper.

Putnam, Robert D., with Robert Leonardi and Raffaella Y. Nanetti. *Making Democracy Work*. Princeton: Princeton University Press, 1993.

Rempel, John K., and John G. Holmes. 1986. "How Do I Trust Thee?" *Psychology Today*, 20(2): 28-34.

Rotter, Julian B. 1967. "A New Scale for the Measurement of Interpersonal Trust." *Journal of Personality*, 35(4): 651-665.

Scharlemann, Jorn, P. W., Catherine C. Eckel, Alex Kacelnik, and Rick K. Wilson. 2001. "The Value of a Smile: Game Theory with a Human Face." *Journal of Economic Psychology*, 22(5): 617-640.

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

Tungodden, Bertil, Alexander Cappelen, Astri D. Hole and Erik Ø. Sørensen. 2007. "The Pluralism of Fairness Ideals?" *American Economic Review*, 97(3): 818-827.

Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science*, 185: 1124-1131.

Wrightsman L. S. 1991. "Interpersonal Trust and Attitudes Toward Human Nature." In *Measures of Personality and Psychological Attitudes* ed. John P. Robinson, Phillip R. Shaver, and Lawrence S. Wrightsman, 373-412. San Diego, CA: Academic Press.

Yaari, M. E. and M. Bar-Hillel. 1984. "On Dividing Justly." *Social Choice and Welfare*, 1(1): 1-24.

**Table 1. Factors influencing individual-level distrust (own demographics and past experiences) (OLS)**

*Dependent variable: average maximum buy-out price (distrust)*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<b>Own Demographics</b>							
Age (above 24)	-0.51 (0.34)			-0.82** (0.32)	-0.72** (0.31)	-0.72** (0.29)	-0.75** (0.29)
Female	-0.75*** (0.21)			-0.45** (0.20)	-0.52** (0.20)	-0.49** (0.19)	-0.55*** (0.19)
Race (Black)	0.33 (0.37)			0.20 (0.32)	0.20 (0.33)		
Race (non-White & non-Black)	-0.92*** (0.29)			-0.72*** (0.26)	-0.81*** (0.26)	-0.85*** (0.24)	-0.88*** (0.24)
Education (college degree)	0.77** (0.32)			0.84*** (0.29)	0.79*** (0.29)	0.82*** (0.26)	0.80*** (0.26)
SES (Lower & Lower-Middle)	0.02 (0.25)			-0.17 (0.23)	-0.20 (0.23)		
SES (Upper & Upper-Middle)	0.08 (0.25)			0.01 (0.22)	0.03 (0.22)		
Only child	0.08 (0.34)			-0.27 (0.31)	-0.18 (0.31)		
<b>Past Experiences and Behaviors</b>							
Leave door unlocked		0.03 (0.11)		0.04 (0.10)			
Lend money		0.30** (0.14)		0.22* (0.13)		0.16 (0.11)	
Lend possessions		-0.07 (0.12)		-0.12 (0.12)			
Trusting behaviors index			0.17 (0.14)		0.03 (0.13)		
Benefit from stranger		-0.74*** (0.22)	-0.79*** (0.22)	-0.80*** (0.20)	-0.82*** (0.21)	-0.77*** (0.19)	-0.76*** (0.19)
Constant	3.02*** (0.21)	2.82*** (0.58)	3.16*** (0.18)	3.32*** (0.55)	3.59*** (0.24)	3.35*** (0.23)	3.55*** (0.19)
R <sup>2</sup>	0.33	0.24	0.19	0.53	0.49	0.50	0.48
Obs.	60	60	60	60	60	60	60

Standard errors in parentheses

\* - p < 0.1, \*\* - p < 0.05, \*\*\* - p < 0.01

**Table 2. Factors influencing individual-level distrust (attitudes) (OLS)***Dependent variable: average maximum buy-out price (distrust)*

	(1)	(2)	(3)	(4)
<b>Own Demographics</b>				
Age (above 24)	-0.51 (0.31)	-0.32 (0.32)	-0.53 (0.33)	-0.39 (0.34)
Female	-0.62*** (0.21)	-0.58*** (0.21)	-0.70*** (0.21)	-0.55** (0.21)
Race (non-White & non-Black)	-0.95*** (0.26)	-0.78*** (0.28)	-0.92*** (0.27)	-0.84*** (0.28)
Education (college degree)	0.79*** (0.29)	0.60** (0.29)	0.75** (0.31)	0.69** (0.31)
<b>Attitudes / Behaviors</b>				
Overall trust attitude (standardized)	-0.19* (0.10)			-0.13 (0.11)
Altruism (standardized)		-0.20* (0.10)		-0.14 (0.12)
Risk attitude (standardized)			0.05 (0.10)	0.01 (0.10)
Constant	3.01*** (0.17)	2.97*** (0.17)	3.07*** (0.17)	2.95*** (0.18)
R <sup>2</sup>	0.36	0.36	0.32	0.38
Obs.	60	60	60	60

Standard errors in parentheses

\* - p &lt; 0.1, \*\* - p &lt; 0.05, \*\*\* - p &lt; 0.01

**Table 3. Distrust, altruism, attitudes, and untrustworthiness by location**

<b>Target Location</b>	<i>Player 1 (laboratory subject)</i>					<i>Player 2 (field participant)</i>
	Avg. buy-out price (distrust)	Avg. amount shared (altruism)	Average attitudinal response (trust) (standardized)			Avg. excess amount kept (untrustworthiness)
			Specific target/context	Related target/context	General trust-worthiness	
(a) U of Pitt.	\$2.75	\$1.64	0.13	0.50	0.48	\$1.89
(b) CMU	\$2.61	\$1.23	0.31	0.58	0.53	\$1.90
(c) Shadyside	\$2.35	\$1.08	0.28	0.27	0.23	\$1.28
(d) Southside	\$2.54	\$1.48	-0.16	-0.35	-0.34	\$1.62
(e) Hill District	\$2.80	\$2.29	-0.55	-1.00	-0.90	\$3.38
<b>Wilcoxon</b>	<i>matched-pairs sign-rank</i>					<i>rank-sum</i>
p<0.10	a-e		a-c	a-b	a-c, b-c	
p<0.05	a-d, b-e	b-d	a-b			b-c
p<0.01	c-d	a-c, a-e, c-d	a-d	a-c, b-c		
p<0.001	a-c, b-c, c-e, d-e	a-b, b-e, c-e, d-e	a-e, b-d, b-e, c-d, c-e, d-e	a-d, a-e, b-d, b-e, c-d, c-e, d-e	a-d, a-e, b-d, b-e, c-d, c-e, d-e	a-e, b-e, c-e, d-e

**Table 4. Factors influencing distrust towards targets (OLS)**

*Dependent variable: maximum buy-out price (distrust)*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<b>Target Demographics: Predicted</b>							
Pred. Target Age (above 24)	-0.18** (0.09)	-0.27*** (0.10)	-0.31*** (0.08)		-0.30*** (0.08)		-0.31*** (0.08)
Pred. Target Gender (Female)	-0.11 (0.08)	-0.13 (0.09)					
Pred. Target Race (Black)	0.04 (0.14)	-0.03 (0.15)					
Pred. Target Race (~White & ~Black)	-0.03 (0.10)	-0.11 (0.12)					
Pred. Target Educ. (no college)	-0.16 (0.16)	-0.10 (0.16)					
Pred. Target Educ. (college degree)	-0.12 (0.10)	-0.09 (0.11)					
Pred. Target SES (Low & Low-Mid)	0.31* (0.16)	0.37** (0.16)	0.40*** (0.08)		0.42*** (0.09)		0.39*** (0.08)
Pred. Target SES (Upp & Upp-Mid)	-0.13 (0.08)	-0.09 (0.09)					
Predicted Same Age		-0.17* (0.09)	-0.12 (0.08)		-0.12 (0.08)		-0.13 (0.08)
Predicted Same Gender		-0.01 (0.09)					
Predicted Same Race		-0.10 (0.10)					
Predicted Same Education		0.08 (0.10)					
Predicted Same SES		0.18** (0.08)	0.19*** (0.07)		0.19*** (0.07)		0.20*** (0.07)
Familiarity Index (standardized)				-0.02 (0.03)	0.02 (0.04)		
Altruism (standardized)						0.04 (0.06)	0.03 (0.06)
Constant	2.76*** (0.07)	2.82*** (0.14)	2.63*** (0.08)	2.61*** (0.03)	2.62*** (0.08)	2.61*** (0.03)	2.63*** (0.08)
R <sup>2</sup>	0.78	0.79	0.78	0.75	0.78	0.75	0.78
Obs.	299 (60)	299 (60)	299 (60)	300 (60)	299 (60)	300 (60)	299 (60)

Subject fixed effects

Standard errors in parentheses

\* - p < 0.1, \*\* - p < 0.05, \*\*\* - p < 0.01

**Table 5. Influence of attitudinal trust on behavioral distrust towards targets (OLS)**

*Dependent variable: maximum buy-out price (distrust)*

	(1)	(2)	(3)	(4)	(5)	(6)
<b>Target Demographics: Predicted</b>						
Pred. Target Age (above 24)				-0.35 <sup>***</sup> (0.08)	-0.35 <sup>***</sup> (0.09)	-0.33 <sup>***</sup> (0.08)
Pred. Target SES (Low & Low-Mid)				0.25 <sup>***</sup> (0.09)	0.28 <sup>***</sup> (0.10)	0.33 <sup>***</sup> (0.09)
Predicted Same Age				-0.17 <sup>**</sup> (0.08)	-0.13 (0.08)	-0.13 (0.08)
Predicted Same SES				0.17 <sup>**</sup> (0.07)	0.19 <sup>***</sup> (0.07)	0.19 <sup>**</sup> (0.07)
<b>Trust Attitudes</b>						
Specific Target / Context (std)	-0.23 <sup>***</sup> (0.05)			-0.21 <sup>***</sup> (0.06)		
Related Target / Context (std)		-0.12 <sup>***</sup> (0.04)			-0.11 <sup>**</sup> (0.05)	
General Trustworthiness (std)			-0.10 <sup>***</sup> (0.04)			-0.07 (0.04)
Constant	2.61 <sup>***</sup> (0.03)	2.61 <sup>***</sup> (0.03)	2.61 <sup>***</sup> (0.03)	2.71 <sup>***</sup> (0.08)	2.69 <sup>***</sup> (0.08)	2.66 <sup>***</sup> (0.08)
R <sup>2</sup>	0.77	0.76	0.76	0.79	0.78	0.78
Obs.	299 (60)	299 (60)	300 (60)	298 (60)	298 (60)	299 (60)

Subject fixed effects

Standard errors in parentheses

\* - p < 0.1, \*\* - p < 0.05, \*\*\* - p < 0.01

**Table 6. Relationship between distrust towards targets and target untrustworthiness (OLS)**

*Dependent variable: maximum buy-out price (distrust)*

	(1)	(2)	(3)	(4)	(5)
<b>Target Demographics: Predicted</b>					
Pred. Target Age (above 24)		-0.27*** (0.09)	-0.32*** (0.08)	-0.32*** (0.09)	-0.29*** (0.09)
Pred. Target SES (Low & Low-Mid)		0.22* (0.13)	0.10 (0.13)	0.13 (0.14)	0.15 (0.14)
Predicted Same Age		-0.13 (0.08)	-0.17** (0.08)	-0.13 (0.08)	-0.14 (0.08)
Predicted Same SES		0.19*** (0.07)	0.17** (0.07)	0.19*** (0.07)	0.19*** (0.07)
<b>Target Untrustworthiness</b>					
Avg. Excess Amount Kept at Location	0.18*** (0.04)	0.12* (0.07)	0.10 (0.07)	0.11 (0.07)	0.12* (0.07)
<b>Trust Attitudes</b>					
Specific Target / Context (std)			-0.20*** (0.06)		
Related Target / Context (std)				-0.10* (0.05)	
General Trustworthiness (std)					-0.07 (0.04)
Constant	2.25*** (0.09)	2.41*** (0.14)	2.54*** (0.15)	2.49*** (0.15)	2.46*** (0.15)
R <sup>2</sup>	0.77	0.78	0.79	0.79	0.78
Obs.	300 (60)	299 (60)	298 (60)	298 (60)	299 (60)

Subject fixed effects

Standard errors in parentheses

\* - p < 0.1, \*\* - p < 0.05, \*\*\* - p < 0.01

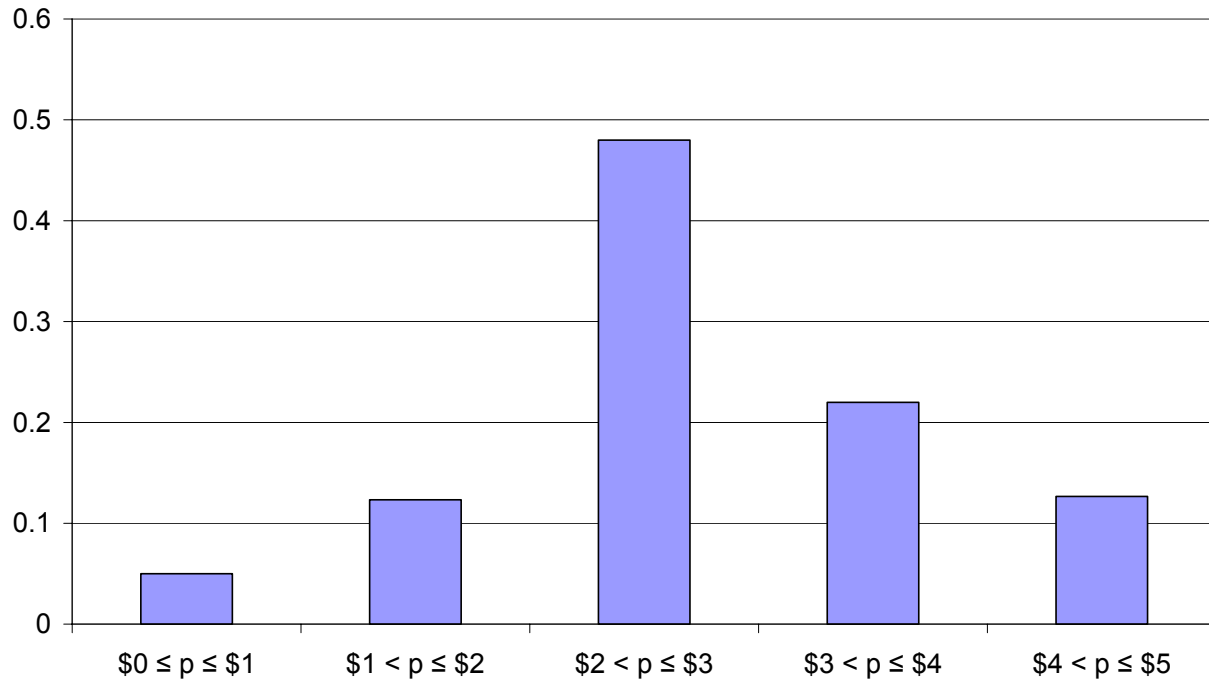
**Table 7. Factors influencing target untrustworthiness (OLS)***Dependent variable: average excess amount kept by target (untrustworthiness)*

	Full sample		Excl. Hill
	(1)	(2)	(3)
<b>Target Demographics: Self-Reported</b>			
Target Age (above 24)	-0.55* (0.28)	-0.66** (0.30)	-0.65** (0.31)
Target Gender (Female)	-0.53** (0.26)	-0.57** (0.26)	-0.70** (0.28)
Target Race (Black)	1.41*** (0.32)	0.65 (0.42)	0.62 (0.43)
Target Race (~White & ~Black)	0.33 (0.38)	0.26 (0.40)	0.17 (0.38)
Target Educ. (no college)	0.64* (0.34)	0.54 (0.35)	0.14 (0.39)
Target Educ. (college degree)	0.14 (0.33)	0.14 (0.33)	-0.07 (0.35)
Target SES (Low & Low-Mid)	0.33 (0.32)	0.22 (0.32)	-0.03 (0.37)
Target SES (Upp & Upp-Mid)	0.84*** (0.32)	0.71** (0.32)	0.54 (0.35)
<b>Locations</b>			
CMU		-0.08 (0.41)	
Shadyside		-0.15 (0.42)	
Southside		-0.15 (0.43)	
Hill District		1.21** (0.53)	
Constant	1.59*** (0.29)	1.82*** (0.37)	2.05*** (0.32)
R <sup>2</sup>	0.13	0.15	0.07
Obs.	300	300	240

Standard errors in parentheses

\* - p &lt; 0.1, \*\* - p &lt; 0.05, \*\*\* - p &lt; 0.01

**Figure 1. Distribution of buy-out prices (N=300)**



**Figure 2. Distribution of average (individual-level) buy-out prices (N=60)**

