

# The Right Amount of Trust\*

Jeffrey Butler  
EIEF

Paola Giuliano  
UCLA, NBER, CEPR and IZA

Luigi Guiso  
European University Institute, EIEF and CEPR

20th April 2010

## Abstract

We investigate the relationship between individual trust and *individual* economic performance. We find that individual income is hump-shaped in a measure of intensity of trust beliefs. Heterogeneity of trust beliefs in the population, coupled with the tendency of individuals to extrapolate beliefs about others from their own levels of trustworthiness, could generate this non-monotonic relationship: highly trustworthy individuals tend to form overly optimistic beliefs, to assume too much social risk and to be cheated more often, ultimately performing less well than those with a belief close to the mean trustworthiness of the population. On the other hand, less trustworthy individuals form overly pessimistic beliefs and avoid being cheated, but give up profitable opportunities, therefore underperforming. The cost of either too much or too little trust is comparable to the income lost by foregoing college. Our findings in large-scale survey data are supported and extended with experimental findings. We show that in the trust game, own trustworthiness and beliefs about others' trustworthiness are strongly correlated and persistent and that patterns in earnings lost due to incorrect beliefs are comparable to those in the survey data.

**JEL Classification:** A1, A12, D1, O15, Z1

**Keywords:** Trust, trustworthiness, economic performance, culture, false consensus

---

\*We thank Oriana Bandiera, Sascha Becker, Larry Blume, Francesco Caselli, Bhajan Grewal, Erzo Luttmer, Bernard Salaniè, Guido Tabellini and Georg Weizacker for helpful suggestions and discussion, and seminar participants at the Bank of Spain, the Einaudi Institute for Economics and Finance, the Kaler Meeting at UCLA, the FEEM conference on the Economics of Culture, Institutions and Crime, the High School of Economics in Moscow, the London School of Economics, University of California Davis, the NBER Political Economy Fall meeting, University of Mannheim, Universidad Pompeu Fabra, University of Siena, Toulouse University for helpful comments. We are grateful to the International Atlantic Economic Society for inviting Luigi Guiso to deliver the Robert A. Mundell distinguished address at the 67th meeting which relied on the material in this paper. Luigi Guiso thanks the EUI for financial support and LUISS University for making the LUISS lab available. Paola Giuliano thanks the UCLA-CIBER grant for financial support.

# 1 Introduction

More than 35 years ago Kenneth Arrow (1972), recognizing the pervasiveness of mutual trust in commercial and non-commercial transactions, went so far as to state that “it can be plausibly argued that much of the economic backwardness in the world can be explained by the lack of mutual confidence” (p. 357). Since then, Arrow’s conjecture has received considerable empirical support. A vast literature investigates the link between aggregate trust and aggregate economic performance and finds a positive and monotonic relationship.<sup>1</sup> However, there is no research available on the relationship between *individuals’* levels of trust—beliefs held about others’ trustworthiness—and *individuals’* economic outcomes. The latter relationship is the focus of this paper.

Trust beliefs are very heterogeneous across individuals. Figure 1 shows the distribution of trust for each of the countries surveyed in the European Social Survey. Here, trust is the belief about the trustworthiness of a generic fellow countryman measured on a scale between 0 and 10: zero means no trust at all, while 10 means other people can be fully trusted.<sup>2</sup> Since for each country beliefs refer to the average trustworthiness of the same population, respondents cannot all be simultaneously right: some must have overly pessimistic beliefs, while others must have beliefs that are too optimistic.<sup>3</sup> Individuals with beliefs in the tails of the trust distribution must either underestimate or overestimate the trustworthiness of others and this should be reflected in their economic performances: those who trust too little will give up trade and profit opportunities too often, depressing their economic performance; conversely, individuals who trust too much will over-invest in others and get cheated more frequently, hampering their economic outcomes. Hence, at the *individual* level, the relationship between trust and economic performance is non-monotonic. There exists an intermediate level of trust—the “right amount of trust”—that maximizes individual income. This amount of income, and trust, will be attained by individuals whose beliefs are closest

---

<sup>1</sup>Trust has been shown to be strongly correlated with GDP per capita and GDP growth (Knack and Zak (2001); Knack and Keefer (1996); Guiso et al (2004); Tabellini (2008b); Algan and Cahuc (2010)); with the ability of firms to grow large (La Porta et. al. (1997)); with the size of a country’s stock market (Guiso et. al. (2008a)); and with cross-country trading patterns (Guiso et al., 2009)).

<sup>2</sup>See Section 4 for the exact wording of the question in the European Social Survey.

<sup>3</sup>We acknowledge that some of the observed heterogeneity in trust beliefs may come from respondents interpreting the question as asking about the trustworthiness of the pool of people with whom they generally interact, rather than asking about the general population. We show in Section 8 that a similar dispersion of trust beliefs exists even without this potential confound in a laboratory experiment. There, participants estimated the trustworthiness of a fixed and explicitly specified population—other experimental participants—and beliefs still exhibited considerable heterogeneity.

to the average trustworthiness in the population.

We test the relationship between trust and income using data from the European Social Survey. Since the survey measures the *intensity* of individuals' trust beliefs on a scale from zero to 10, we can explore the relationship between individual trust and individual economic performance, particularly at the tails of the distribution of trust beliefs.

When we regress individuals' income on a set of dummies for the 11 different levels of trust we find a marked hump-shaped relationship: people with low levels of trust have significantly lower income than those with intermediate levels of trust. Income tends to reach a peak at a level of trust around 8, before declining rapidly for the highest levels of trust.

The magnitude of this income variation is economically significant. On average, income for individuals with the lowest level of trust is 14.5% lower than income at the right amount of trust. This shortfall in income is on par with the income lost by foregoing college. At the highest level of trust, average income falls short of peak income by 7%. Accordingly, the cost of deviating from the right amount of trust can be substantial.

To strengthen the causal interpretation we re-estimate the relationship between individual income and trust beliefs for low- and high-trust countries separately. In countries with low average trustworthiness the right amount of trust is lower than in countries with high average trustworthiness. Hence, the income-trust relationship should, *ceteris paribus*, peak at lower levels of individual trust in the first group of countries than in the second.<sup>4</sup> Consistent with this implication we find that the income-trust relationship is always hump-shaped, but in low trust countries peaks at a lower trust level. We also show that the hump-shaped income-trust relationship is not a reflection of obvious forms of unobserved heterogeneity. For example, we can rule out the explanation that moderate levels of trust actually serve as a generic measure of moderate traits, which are better suited to achieve economic success than extreme traits. Furthermore, we show that the hump-shaped income-trust relationship does not vanish with experience, nor with education.

The European Social Survey allows us to dig deeper into one of the mechanisms through which generalized trust affects economic performance: exposure to the risk of being cheated. The survey asks individuals whether, over the past five years, they have been "cheated"

---

<sup>4</sup>We assume that trust beliefs are, on average, correct so that, e.g., highly-trusting countries are also highly-trustworthy countries.

over four different domains: dealing with a bank, buying goods second hand, buying food, and dealing with a plumber, builder, mechanic or repairman. All else equal, exceedingly trusting individuals should be cheated more often. Obscuring this relationship in the data, however, is the fact that individuals who are cheated learn and revise their trust beliefs downward. In this way, learning generates a negative correlation between trust and the experience of being cheated. To isolate the causal effect of trust on the probability of being cheated we follow two avenues. First, we adopt an instrumental variables approach and find that those who trust more are indeed more likely to be cheated across all the domains for which we have data.

The second avenue we follow involves detailed information on immigrants from the European Social Survey. One strand of literature has shown that cultural norms such as individual trustworthiness and attitudes concerning trusting others are acquired through intergenerational transmission (Giuliano (2007); Fernandez and Fogli (2009); Guiso et. al. (2004, 2008b), Tabellini (2008a, 2008b), Luttmer and Singhal (2010), Dohmen et. al. (2007)) and thus persist across generations. This provides an alternative source of exogenous variation which we use to identify the causal effect of trust on exposure to social risk. Under the hypothesis that trust beliefs persist—because, e.g., cultural norms are slow to change and individuals extrapolate from their own type—immigrants from high trust countries should be, all else equal, more likely to be cheated. Consistent with this prediction, we find that immigrants from relatively high trust countries are, *ceteris paribus*, cheated more often. For first generation immigrants, we find a strong positive effect of average trust levels in the country of origin on being cheated. The second generation learns from experience and, consequently, the impact of trust levels in the country of origin on immigrants' exposure to fraud vanishes.

Finally, to provide more controlled evidence on both the hump-shaped relationship between trust and economic performance and the heterogeneity of trust beliefs we report evidence from an experiment using the trust game (Berg, Dickhaut and McCabe, 1995). We show that own trustworthiness and expectations of others' trustworthiness are strongly correlated and that this correlation persists over multiple rounds of game-play. We also find that participants' trustworthiness, as measured by their behavior in the trust game, can be traced back to values instilled by their parents. Finally, using a modified trust game more closely matching our theoretical model we show that individuals who depart in either dir-

ection from the right amount of trust earn less, and that in percentage terms their earnings shortfalls are of the same order of magnitude as in the survey data.

Our paper contributes to two recent literatures that study how trust relates to various dimensions of economic performance using either field data (see among others, Tabellini (2008b); Algan and Cahuc (2010); Nunn and Wantchekon (2009); La Porta et. al. (1997); Bloom et al (2009)) or laboratory experiments (see Camerer (2003) and the references therein). However, while both of these literatures concentrate on how trust (or lack of it) affects collective economic performance, we combine both methods to study how differences in individuals' trust beliefs result in differences in individuals' economic outcomes.

The remainder of the paper is organized as follows. In Section 2 we discuss alternative explanations for why trust beliefs may differ and persist. In Section 3 we present a simple model that predicts a hump-shaped relationship between individual trust and performance. In Section 4 we describe the survey data and in Section 5 we present the main results from our estimation of the trust-performance relationship in this data. Various extensions and robustness checks follow in Section 6. In Section 7 we estimate the effect of trust on the frequency with which one is cheated. In section 8 we present the results from our trust game experiments. Section 9 concludes.

## **2 Why are trust beliefs heterogeneous and persistent?**

There are two plausible explanations for why trust beliefs differ across individuals in the same community and why these differences persist. According to one view, individuals' beliefs are initially acquired through cultural transmission and then slowly updated through experience from one generation to the next. Heterogeneity is then the result of family-specific shocks. This line of argument has been pursued by Guiso, Sapienza and Zingales (2008b) who build an overlapping-generations model in which children absorb their trust priors from their parents and then, after experiencing the real world, transmit their (updated) beliefs to their own children. Dohmen et. al (2007) provide evidence consistent with this view. This explains heterogeneity in initial priors and persistence across generations. Within a generation, correlation between current beliefs and received priors is diluted as people age and learn. Yet this dilution needs not to be complete. One mechanism generating persistence could be confirmation bias—a tendency to seek and find evidence that

confirms existing beliefs and ignore or reinterpret disconfirmatory evidence.<sup>5</sup> Alternatively, cultural beliefs may persist because, once hardwired, they are painful to eradicate and this pain makes one reluctant to update them even in spite of disconfirmatory information.<sup>6</sup>

The second plausible explanation is that parents instill values rather than beliefs in their children. In particular, parents may teach values of trustworthiness: acting justly even at the expense of material gains. Cultural transmission of values of cooperation and trustworthiness is the focus of Bisin and Verdier (2000), Bisin, Topa, and Verdier (2004) and Tabellini (2008a). They show how norms of behavior are optimally passed down from parents to children and persist from generation to generation. Heterogeneity in parents' preferences and experiences may then result in heterogeneity in instilled trustworthiness. Even if parents do not teach beliefs directly, individuals may extrapolate from their own types when forming beliefs about others' trustworthiness. As Thomas Schelling once wrote "...you can sit in your armchair and try to predict how people behave by asking yourself how you would behave if you had your wits about you. You get free of charge a lot of vicarious empirical behavior" (1966, p. 150).

Through such "false consensus" (Ross, Green and House (1977)), heterogeneity in values translates into heterogeneity in beliefs. In our context false consensus implies that highly trustworthy individuals will tend to think that others are like them and form overly optimistic trust beliefs, while highly untrustworthy people will extrapolate from their own types and form excessively pessimistic beliefs. If values persist over time and false consensus does not vanish with learning, then wrong beliefs will also persist.<sup>7</sup> Both highly trustworthy and highly untrustworthy individuals will tend to systematically form more extreme trust beliefs than are warranted by their experiences.

In the next section we present a simple static model illustrating how heterogeneous trust beliefs can imply a humped-shaped relationship between income and trust.

---

<sup>5</sup>A recent functional magnetic resonance imaging study shows where in the brain the confirmation bias arises and how it operates unconsciously (Westen, et. al., 2006).

<sup>6</sup>Blanco (2008) provides evidence consistent with this view in the context of a trust game experiment. She shows that senders' trust decisions respond to their own values even when the behaviour of the receiver is fully known.

<sup>7</sup>False consensus has been shown to be a persistent phenomenon: neither providing additional information about the population of interest, nor warning individuals about the possibility of false consensus, eliminates the effect (Krueger and Clement (1994)). Furthermore, it has been found that false consensus is not drowned out by monetary incentives for accurate predictions (e.g. Massey and Thaler (2006)).

### 3 Individual trust and economic performance: a simple model

Consider an investor with an endowment  $E$  which can be invested, totally or partially, in a venture managed by a partner. The endowment and the partnership should be interpreted broadly. The endowment could represent capital contributed to a project run by an entrepreneur or money invested in a fund managed by a professional which affects income from capital. Alternatively,  $E$  could be the time and effort (human capital) that a worker allows his or her boss to manage in hopes of advancing more quickly along the career path which affects labor income. Or, in a more familiar setting,  $E$  could represent ideas that a researcher shares with co-authors on a joint project.

An amount  $S \leq E$  invested creates surplus according to the production function  $f(S) > S$ , of which the partner agrees to return a fraction  $0 < \gamma < 1$  to the investor. Partners can be one of two types: honest or cheater. A fraction  $\pi$  of partners are cheaters, while the rest of the economy's partners are honest. Each investor is randomly matched with a partner, as in Dixit (2003). An honest partner returns the promised share of the surplus,  $\gamma f(S)$ , while a cheater absconds with the whole surplus. We assume that  $f(S)$  is increasing and concave ( $f' > 0$  and  $f'' < 0$ ), and that  $\gamma f(S) > S$  so that investment has a positive return if the partner does not cheat. We also assume that  $(1 - \pi)\gamma f'(0) > 1$ : at zero investment the expected marginal return from investing exceeds the return from keeping the endowment idle. Together these assumptions imply a unique, internal, optimal investment amount.

Investors also differ in their intrinsic trustworthiness. Assume there is a continuum of investors each characterized by a level of trustworthiness,  $p$ , distributed on the unit interval  $[0, 1]$ . Suppose that a fraction  $\pi$  of investors has a level of trustworthiness  $p \leq \pi$ . Of course, while an investor can be cheated, investors cannot cheat. We first assume that investors have correct beliefs and anticipate that the probability of being cheated is  $\pi$ . Given these correct beliefs, an individual investor decides how much to invest in the venture so as to maximize:

$$\begin{aligned} \text{Max}_S : Y(S) &= E - S + (1 - \pi)\gamma f(S) \\ \text{s.t.} : S &\leq E \end{aligned}$$

Let  $S_\pi^*$  denote the optimal amount invested when beliefs are correct. Then the investor's expected income is  $Y(S_\pi^*) = E - S_\pi^* + (1 - \pi)\gamma f(S_\pi^*)$ .

Consider next the case where investors form beliefs about the trustworthiness of the partners by extrapolating from their own types. To illustrate, suppose that the trust belief of an individual with own-trustworthiness level  $p$  is exactly his or her own trustworthiness. Given these (possibly incorrect) beliefs, an investor solves the problem:

$$\begin{aligned} \text{Max}_S : Y(S) &= E - S + (1 - p)\gamma f(S) \\ \text{s.t.} : S &\leq E \end{aligned}$$

Let  $S_p^*$  denote the optimal amount invested when beliefs about others' trustworthiness are extrapolated from the investor's own type and let  $Y(S_p^*) = E - S_p^* + (1 - \pi)\gamma f(S_p^*)$  be the investor's average income. Notice that income realizations depend on the actual fraction of trustworthy partners. We state the following proposition:

**PROPOSITION 1:** When individuals extrapolate from their own type, an investor's average income,  $Y(S_p^*)$ , is a concave function of the investor's trust beliefs. This function attains its maximum when the investor's belief about the share of trustworthy partners,  $1 - p$ , equals the true share  $(1 - \pi)$ . Proof: see appendix 1.

The proposition implies that both investors with very low and very high levels of trust (and trustworthiness) do worse than those with a trust (and trustworthiness) level close to the average trustworthiness of the population. In the first case, by under-investing investors with very low trust levels lose little if cheated; but by retaining too much of their endowment, they give up profit opportunities—and the latter effect far exceeds the former. On the other hand, investors in the second group invest a lot in the productive venture, which can potentially raise their income. But since they grant partners more trust than they deserve, they lose a lot when cheated and the latter effect dominates the former. Hence, the relationship between individual economic performance and trust beliefs is hump-shaped, as illustrated in Figure 2.

If populations differ in their average degree of trustworthiness  $(1 - \pi)$ , then observed individual performance,  $Y(S_p^*)$ , will, *ceteris paribus*, be higher in high-trustworthiness countries. In our empirical analysis, we will capture these shifts with a set of country dummies as well as finer, community-level, trustworthiness controls.

Furthermore, our model makes a prediction about how the income-trust relationship will vary across countries:



**PROPOSITION 2:** The level of trust at which income attains its maximum is increasing in the population’s proportion of trustworthy people,  $(1 - \pi)$ .

This is a relevant prediction that can be tested empirically.<sup>8</sup> With respect to Figure 2, the proposition states that a larger share of trustworthy individuals implies a peak in income that occurs farther to the right.

## 4 Data: the European Social Survey

To study the relationship between individual performance and trust beliefs we rely on the second wave of the European Social Survey (ESS), conducted in 2004/2005. The ESS is a biennial cross-sectional survey administered in a large sample of mostly European Nations. The survey has been conducted four times: in 2002/2003, 2004/2005, 2006/2007 and 2008/2009. The number of countries varies by wave.<sup>9</sup> We use the second round because it is the only round containing the measures of cheating crucial for our analysis. For each country, the ESS provides information on individuals’ social values, cultural norms and behavior patterns. Within each country, a representative sample of around 2,000 individuals is surveyed.<sup>10</sup> Pooling data across countries yields 47,035 observations. Because some observations lack income data (11,128 observations) or responses to one or more explanatory variables (7,624 observations) our final reference sample consists of 28,383 observations. The data appendix provides details on the sample selection, the countries included and the

---

<sup>8</sup>There are two points to notice. First, in this simple model the channel through which trust beliefs and individual performance are related is systematic errors in beliefs. We argue that these can be induced by false consensus or by ingrained heterogeneous priors. Obviously, there could be other channels. For instance, high-trust people may become targets of swindlers who can exploit naïve expectations of good faith. Alternatively, highly-trusting people may be more exposed to confidence games even when their own attitude to trust is not explicitly targeted. Barnard Madoff’s story can be interpreted as one where high-trust individuals were more likely to fall prey to Madoff’s game even if they were not individually targeted. These two mechanisms can explain why those who trust too much may under-perform. However, they cannot explain why those who trust too little do poorly. Culturally-induced heterogeneity in beliefs or in values, together with the tendency of individuals to extrapolate from their own types, can explain both. Second, the model implies that in the absence of false consensus all individuals would share the same beliefs even if individuals are actually heterogeneous. False consensus, by linking trust beliefs to own-trustworthiness, automatically gives rise to heterogeneity in beliefs. In this context, one interpretation of false consensus is to view it as a source of initial priors, which allows for a departure from the controversial common prior assumption. In the absence of a history of information about the reliability of a pool of people, those interacting with an unknown pool form a prior by asking themselves how they would behave in similar circumstances: since they would personally behave differently, they start with different priors. This is consistent with the evidence shown in Section 7.

<sup>9</sup>There are 22 countries included in the first round, 26 in the second, 25 in the third and 28 in the latest.

<sup>10</sup>Sample size differs by country depending on country population and ranges from 579 in Iceland to 2870 in Germany.

overall survey design. Besides information on core variables of interest for the purpose of the survey, the ESS provides various measures of respondents’ demographic characteristics and indicators of respondents’ socioeconomic status.

#### 4.1 Measuring trust

The ESS elicits trust beliefs by asking the classical question “Generally speaking, would you say that most people can be trusted, or that you can’t be too careful in dealing with people?” While in most comparable surveys (the World Values Survey, the US General Social Survey, etc.) the trust question is binary, in the ESS respondents are asked to express the *intensity* of their trust beliefs on a scale from 0 to 10, where 0 means no trust at all and 10 means that most people can be fully trusted.<sup>11</sup> It is this feature of the ESS that allows us to investigate the relationship between individuals’ trust beliefs and economic performance, and, in particular, to test whether the relationship is hump-shaped.

Figure 1 illustrates both the presence of a considerable number of observations at the two tails of the distribution of trust within each country, as well as systematic differences in the shape of the trust distribution across countries. In one group—the high trust North European countries such as Norway, Denmark, Finland, Sweden and the Netherlands—the distribution has a fat tail on the right and the modal level of trust is quite high at around 7 or 8. Another group, which includes the Mediterranean countries and several Eastern European countries, features a fat left tail, denoting low average trust. In a third group consisting of several European countries (e.g., Austria, Germany, France and the UK) the distribution is approximately symmetric around modal values of 5 or 6. Table 1 shows that the (whole) sample mean trust level is around 5, with a standard deviation of 2.5. As we show in Section 5, both within- and between-country variation prove critical for identifying the predicted hump-shaped relationship between performance and trust.<sup>12</sup>

---

<sup>11</sup>In most comparable surveys, respondents must choose between the two options: “most people can be trusted,” or “you can’t be too careful.”

<sup>12</sup>In the final sample used in our estimations, there are over 2,000 observations featuring a trust level of 9 or 10 and almost 4,000 with trust equal to either 1 or 0. We investigated the characteristics of individuals reporting a very high level of trust to address the concern that there could be something peculiar about these people, other than trust, driving their incomes down. We found no systematic differences between these highly-trusting individuals and the rest of the sample. In particular, their education levels were not significantly different.

## 4.2 Measuring performance

The ESS is rich in many dimensions, but as with most surveys focusing on values it has little information on individuals' economic outcomes or other economic variables. The best available performance indicator is a measure of total net household income, which is the measure we use. Respondents are asked to report which income bracket, identified with a letter, best approximates their households' total net income. They are asked to consider income from all sources, including labor income and income from capital and investments. This is an important feature because, as we have argued, trust can affect all sources of income.<sup>13</sup> In order to facilitate answers, the question is framed in a way that accounts for country-specific conventions in the frequency of income payments. Respondents can provide the income figure using the frequency they know best: weekly, monthly or annual. Each letter identifies an income bracket in euros (see appendix 2 for more details) defined so as to be consistent across different frequencies.<sup>14</sup> We convert all responses to their annual equivalent. The resulting brackets range from less than 1800 euros per year to above 120,000 (the largest net income allowed). To facilitate our analysis we identify each bracket with its mid-point. Table 1 shows summary statistics for (log) income in the sample.

## 5 Model specification and results

To study the relationship between individual economic performance and individual trust we estimate the following model:

$$\log(Y_{ic}) = \sum_j \alpha_j Trust_{jic} + \beta X_{ic} + \delta C + \eta R + \epsilon_{ic} \quad (1)$$

Here  $Y_{ic}$  is the income level of individual  $i$  in country  $c$  and  $X_{ic}$  is a vector of individual controls that can affect economic performance. We capture the effect of trust with a set of 10 dummies  $Trust_1, Trust_2, \dots, Trust_{10}$ , the excluded group being individuals reporting the lowest possible trust level of 0. This specification allows wide flexibility estimating the relationship between trust and income, imposing no parametric assumptions. Finally, to

---

<sup>13</sup>While trust may affect all types of income, certain types, such as income from capital, may arguably be more exposed to opportunistic behavior than other types (e.g., labor income) and thus more sensitive to incorrect trust beliefs. Unfortunately, we cannot test this possibility as the ESS does not provide information on income components.

<sup>14</sup>For instance, the first income category identifies income below 40 euros per week or below 150 euros per month or below 1800 euros per year. These figures are roughly equivalent if a month is made of four paid working weeks and a year of 12 paid working months.

control for systematic differences in average income across countries and, within countries, across regions, we insert a vector of country fixed effects,  $C$ , and a vector of regional fixed effects,  $R$ . Among other things, these fixed effects capture differences in individual performance due to systematic differences in the average level of trustworthiness across countries and, within countries, across different regions. The vector  $X$  includes dummies for individuals' educational attainment as well as dummies for respondents' fathers' levels of education, which serve as proxies for acquired and inherited human capital, respectively. The vector  $X$  also contains a linear and quadratic term in age to capture life cycle effects in income, dummies for gender, marital and employment status, immigration status, as well as dummies for city size with rural areas as the excluded category.

Table 2 shows the resulting estimates. We report neither country nor regional fixed effects, instead showing only the  $F$ -test for their joint significance in the note at the bottom of the table. The first column reports estimates for the whole sample, using as regressors the set of trust dummies, other plausible determinants of income that may also be correlated with trust (Alesina and La Ferrara, 2002) as well as country and region fixed effects. The income-trust relationship is increasing for low levels of trust, before leveling off and peaking at a trust level of 8. Beyond a trust level of 8 income declines. The decline is initially small. However, income falls precipitously moving trust level 9 to the highest trust level. This concave pattern is evident in Figure 3 and is fully consistent with our model's predictions (Section 3). Most of the other controls have effects consistent with our priors: income increases with own and father's education; it is hump-shaped in age; it is higher for males and for married or single people (compared to those who are divorced or widow(er)s). Income is lower for the unemployed, for those out of the labor force and for immigrants.

Concerning the magnitude of the effect of trust beliefs on income, those with the lowest level of trust have an income that is 14.5 percent lower than income at the peak, a decrease of the same order of magnitude as the increase (15.7 per cent) associated with obtaining a college degree. Those who express the highest level of trust (10) make an income that is 7.3 percent lower than peak income. Both of these differences are statistically significant, as the  $t$ -tests at the bottom of the table show. Thus, departing from the right amount of trust, either because one trusts too much or because one trusts too little can be individually very costly.

In our estimates we are interpreting measured trust as picking up only individual beliefs

about others’ trustworthiness. There is a still-unsettled debate over whether questions such as those asked by the ESS or the World Values Survey reflect expected trustworthiness only, or reflect a mix of beliefs and individual preferences (see Miller and Mitamura (2003)). For instance, Fehr (2009) points out that answers to trust questions like those asked in the ESS likely reflect not only individuals’ beliefs about others’ trustworthiness, but also individuals’ preferences towards risk, and in particular towards social risk.<sup>15</sup> Alternatively, it has been argued (Cox, 2004) that trust may reflect pure altruistic preferences in addition to beliefs about others’ trustworthiness, so that for given beliefs more altruistic individuals would exhibit more trust. This would be the case, for example, when trust is measured as the act of sending money in standard trust games, but could also be reflected in survey measures of trust if individuals respond by mimicking what they would do in an experiment. In Section 8 we more fully address these issues using evidence from a trust game experiment that allows us to separate beliefs from preferences. Here, to account for these possibilities we add a control for risk preferences (column 2), as well as a proxy for altruism (column 3). Our measures rely on questions eliciting attitudes on various domains by asking participants how a certain description applies to them. Respondents were asked the following question: “I will briefly describe some people. Please listen to each description and tell me how much each person is or is not like you.” To obtain an indicator of risk attitudes we use the following description: “She/he looks for adventures and likes to take risks. She/he wants to have an exciting life.” To obtain an indicator of altruism we rely on the following description: “It is important to her/him to be loyal to her/his friends. She/he wants to devote herself/himself to people close to her/him.”<sup>16</sup>

For these questions, respondents provide answers between 1 and 6, with 1 meaning “very much like me,” 6 meaning “Not like me at all” and values in between reflecting intermediate similarity. Thus higher values of the risk preference indicator signal high risk aversion and higher values of the altruistic preferences measure mean less altruism. In all of the analysis that follows, we re-order responses to these questions so that higher values indicate higher risk tolerance and more altruism, respectively.

---

<sup>15</sup>Glaeser et. al. (2000) argue that WVS questions are better measures of trustworthiness than of trust beliefs as they correlate poorly with amounts sent in a trust game. Sapienza et. al (2007) argue that this is due to behavioral trust measuring both beliefs and preferences. They conduct a trust game where subjects report expected trustworthiness and are also asked the standard WVS generalized trust question. They find that a sender’s expectations of others trustworthiness is highly correlated with the trust question in the WVS, suggesting that the latter is a good measure of the belief component of trust.

<sup>16</sup>Two other measures of altruistic preferences from questions in the survey provide very similar results

Risk tolerance is positively and significantly correlated with income while measures of altruism are in general negatively correlated. However, when we control for these preference measures, and thus net out the trust measure from their influence, our results are unaffected. This suggests that the trust-performance relationship is not merely an artifact of trust serving as a proxy for risk preferences or altruism.

Finally, in the last column we replace the set of trust dummies with a linear and quadratic term in the trust variable. This parametric specification allows a direct test of concavity. Consistent with the previous evidence, the linear term is positive and significant, while the quadratic term is negative and significant. Using the estimated parameters, the maximum level of income is obtained when trust is equal to 7.5, confirming the hump-shaped relationship.

### **5.1 In medio stat virtus: unobserved heterogeneity**

One concern with the estimates in Table 2 is that the trust measure may be capturing a general tendency of individuals with moderate attitudes (e.g., moderate risk aversion or moderate generosity) to succeed economically. For instance, it may be that people who are too generous or too stingy make less income than moderately generous people, and moderation itself is an individual characteristic also reflected in moderate levels of trust. After all, Aristotle a few millennia ago theorized that those who live a balanced life and avoid excess can achieve happiness. This balance, he taught, varies among different persons and situations, and exists as a golden mean between two vices—one an excess and one a deficiency.

In Table 2 we controlled for a variety of traits, but the effect of these traits on income was assumed to be linear. To deal with the possibility that these variables have, themselves, a hump-shaped effect captured by trust, we include in Table 3 a full set of dummies (excluding the lowest category) for each of the traits considered. This allows extreme attitudes to affect income non-monotonically. In addition, we expand the set of traits to include a measure of religiosity on a 0-10 scale (ranging from “not religious at all” to “very religious”) and an index of political ideology (also on a 0-10 scale, from far left to far right) as measures of moderation. In all cases we find that the concavity of income in trust is statistically robust to this re-specification. Furthermore, except for political opinions, income is not hump-shaped in these traits.

## 5.2 Reverse Causality

Another, perhaps more serious, concern when looking at the correlation between individual income and trust is that it may be income causing patterns in trust rather than the other way around, as we are arguing. For instance, high income people may be more prone to trust others if they tend to accumulate more social relations, as in Glaeser (2000), and social relations enhance trust. Insofar as this reverse causality argument is true, the rising portion of the trust-performance relationship we document may reflect it; however it cannot explain the declining part of the relationship. Similarly, if for whatever reason high income causes lower trust, then reverse causality could explain the falling part of the relationship but not its rising portion. Hence reverse causality, even if present, cannot be the full driver of the relationship; on the other hand, culture-driven diversity in trust beliefs is able to explain both the rising and the falling parts. Indeed, identification of the effect of trust on individual performance occurs through the non-linearity of the predicted relationship.

To dig deeper into the reverse causality mechanism, we exploit our model’s implication that income should peak at different levels of trust in populations with different average trustworthiness levels (Section 3, Proposition 2). Since we do not directly observe trustworthiness, we obtain a measure of a country’s average level of trustworthiness by taking averages of the levels of trust of the individuals in that country. If people actually extrapolate trust beliefs from their own trustworthiness—as in our model in Section 3—then the average belief is a good measure of average trustworthiness.

In Table 4 we report results of our basic model when we distinguish between below-median, above-median and median trust countries (first two columns), or when we estimate the model separately for countries with a fat tail to the right (the high-trust countries) and a fat tail to the left (the low trust countries) on the basis of Figure 1.<sup>17</sup> As Table 4 and, even more clearly, Figure 4, show, in countries with below-median trust income peaks at a trust level of 7; on the other hand it does not reach its peak until trust equals 9 in above-median trust countries. Consistent with our model’s causal mechanism, trusting a lot can be particularly harmful in countries where the share of untrustworthy people is large. In these countries, fully trusting others reduces income by 10 percentage points compared

---

<sup>17</sup>We define high trust countries those with fat tails to the right (Switzerland, Denmark, Finland, Iceland, Netherland, Sweden and Norway). We define low trust countries those with fat tail to the left (Greece, Italy, Portugal, Turkey, Czech Republic, Hungary, Poland, Slovenia and Slovakia).

to the peak. In contrast, the analogous loss of income is only 4.6 percentage points in high-trust countries.

The evidence so far lends credence to the idea that reverse causality is not the primary driver of our results. Section 8 below—the experimental part of the paper, in which income is exogenously given to subjects—further helps to rule out this possibility.

## 6 Extensions and robustness checks

In this section we extend our results and test their robustness to a variety of specifications and additional controls.<sup>18</sup> First, if the concave relationship between individual trust and performance reflects the tendency to extrapolate beliefs from one’s own type, an obvious question is whether this impacts all groups equally or whether there are some groups that are more prone to it. The literature maintains that false consensus is persistent and universal. Hence, it should not vanish with experience and should not disappear among individuals with high cognitive ability. In our context, these two properties suggest that the trust-performance relationship should retain its humped shape when we group individuals according to experience or ability. Therefore, we re-estimate our model distinguishing first between “young” and “old” (individuals below age 45 and above, respectively) using age as a proxy for experience; and second, between high and low education (more or less than a secondary education, respectively) as a proxy for cognitive ability. Neither of these two distinctions makes a qualitative difference: the trust-performance relationship is as concave among the “young” as it is among the “old,” and equally concave among those with a low level of education as it is among those with a high level of education. This is apparent in Figure 5.<sup>19</sup>

Second, one may argue that the concavity in the relationship between trust and income is the result of systematic variation in the dispersion of trust beliefs with income. To be clear, suppose that individuals can collect costly information about the probability that their counterparts are trustworthy. Wealthier people can afford to pay for more informative signals about their trading partners and therefore have more precise assessments of their trustworthiness. If true, this implies that wealthy people have similar trust beliefs concentrated around the population’s true trustworthiness; the middle class would have beliefs

---

<sup>18</sup>To save space, we report regression estimates from this section in appendix 4.

<sup>19</sup>Regression results are reported in Table A1



that are correct on average but somewhat less precise; while the poor would also have beliefs that are correct on average but even more diffuse. Though no systematic relationship would exist between trust beliefs and income, heterogeneity in belief precision could mechanically imply a hump-shaped relationship between trust and economic performance. This difference in incentives to collect information has, however, another implication: dispersion in trust beliefs should be inversely related to income. To check whether this mechanism is driving our results, we computed for each country the relationship between the standard deviation of trust beliefs and income. The predicted negative relationship is not in the data (Fig. 1A, Appendix 4).

Third, while we have focused on the effect of trust on income, one may wonder whether our estimates are not also reflecting an effect of trustworthiness on income. In our model trustworthiness affects income only because it affects trust beliefs through false consensus. But trustworthiness may pay off directly, e.g., because trustworthy people are delegated more responsibility and this is remunerated. Thus, one concern is that the hump-shaped relationship could be driven by trustworthiness and not trust. To check this possibility we add to our basic income regressions a quadratic specification for several proxies of trustworthiness (Table A2). Trustworthiness is proxied in two ways: i) using self-reported cheating behavior in the past 5 years;<sup>20</sup> and ii) using an index of how much delegation a worker is granted in the workplace.<sup>21</sup> While trustworthiness tends to affect income (monotonically) positively, the relationship between trust beliefs and income remains unaltered.<sup>22</sup>

Finally, to test whether the hump-shaped effect of trust on income is robust to the inclusion of an even larger set of controls than already considered, we ran a set of richer specifications. To our standard regressions we added a full set of education dummies and interactions with each country in the sample, and measures of mothers' and partners' education as well as the number of people living at home. The concavity of the trust-income relationship remained unaffected as did the trust dummy coefficients (Table A3).

---

<sup>20</sup>Specifically: how often the person kept the change from a shop assistant/waiter when given too much; whether he paid cash with no receipt to avoid VAT or tax; sold something second-hand and concealed its fault; misused or altered card/document to pretend to be eligible; made an exaggeration or false insurance claim; offered a favour/bribe to a public official in exchange of his service; falsely claim government benefits)

<sup>21</sup>The variable, called "trustworthiness" in the table, is described in detail in Section 7.2.

<sup>22</sup>Note that three measures of trustworthiness have a negative impact on income. Those measures are the ones indicating whether the respondent kept change from a shop assistant or waiter, or whether they made a false insurance claim or whether they falsely claimed government benefits. This is possibly due to these measures being an indication of honesty rather than trustworthiness, and there being an explicit monetary prize associated with dishonesty.

## 7 Trust and cheating

Two sources of suboptimal behavior contribute to the concavity of income in trust. On one hand, too little trust worsens performance through overly-cautious decision making that leads to missed profit opportunities. On the other hand, too much trust undermines performance by increasing the chances of being cheated and, conditional on being cheated, exposing individuals to larger losses. The first channel implies that the chances of missing profitable opportunities are smaller for those who trust more; the second channel implies that the chances of being cheated are increasing in trust. Providing evidence on the first channel is problematic because missed opportunities are typically unobservable. However, because we have information on how often individuals have been cheated along various domains we can test the second channel.

### 7.1 Measuring cheating experience

The ESS reports information on how often respondents have been cheated within the five years prior to the interview along four dimensions: being cheated by a bank/insurance company; a plumber, builder, car mechanic or other repair person; a seller of second hand goods; or a grocer or food seller. Specifically, the ESS asks participants:

“How often, if ever, have each of these things happened to you in the last five years?”

1. A bank or insurance company failed to offer you the best deal you were entitled to.
2. You were sold something second-hand that quickly proved to be faulty.
3. You were sold food that was packed to conceal the worse bits.
4. A plumber, builder, car mechanic or other repair person overcharged you or did unnecessary work.

Respondents could answer in one of 5 ways—never, once, twice, 3 or 4 times or, finally, 5 times or more—which we code with the numbers 0 to 4. Figure 6 shows histograms of the answers to each of the four cheating dimensions for the pooled data. Not surprisingly, in all cases there is a spike at “Never,” so that the vast majority of respondents report not having been cheated. However, a non-negligible proportion of people—ranging from 22% in the case of the purchase of second-hand goods to over 40% for food purchases—report

having been cheated at least once. Furthermore, quite a few people report being cheated more than once, but the frequency with which people report being cheated decays rapidly in all domains except food, where close to 10% of respondents report being cheated 5 times or more.

In addition to analyzing the frequency with which individuals are cheated in each of the four domains, we also construct two summary indicators: the number of times an individual has been cheated over the four domains collectively, and a variable extracting the first principal component of the four cheating indicators. Summary statistics are reported in Table 1.

## 7.2 Empirical specification

To test whether the chances of being cheated increase with trust we estimate the following model:

$$Z_{ic}^d = \gamma Trust_{ic} + \delta X_{ic} + \lambda C + \theta R + \xi_{ic} \quad (2)$$

where  $Z_{ic}^d$  is an indicator of how often individual  $i$  has been cheated in country  $C$  in the domain  $d$  (cheated by a bank; or when buying food; or by a car repairer; or when buying goods second hand). The other variables have the same meaning as in (2), but in this specification trust is a single variable taking values from 0 to 10.<sup>23</sup> We control for income to capture differences in the number of transactions people engage in for a given level of trust. In addition, to address the concern that trust is simply a proxy for risk attitudes, we add the survey measure of risk tolerance as a control.

Moreover, we insert into this regression a full set of country and region dummies to account for national and intra-national differences in the fraction of cheaters, and to absorb any location-specific characteristics that may encourage or discourage cheating.<sup>24</sup> Heterogeneous cultural priors or heterogeneity in trustworthiness coupled with false consensus generates dispersion in trust beliefs across individuals around these location-specific means, exposing them differentially to the possibility of being cheated. These differential effects are what our regression will be capturing.

Before showing the estimates of (2) we have to confront an identification issue. Since

---

<sup>23</sup>We use a single trust measure rather than a set of dummies because we are going to instrument trust. Furthermore, in principle being cheated should increase monotonically with trust.

<sup>24</sup>These fixed effects also take care of any variation across countries and regions in what is considered to be cheating, and that may result in different frequencies of reported cheating across countries and regions.

people learn from experience and revise their priors accordingly, those who have been cheated are more likely to revise their trust beliefs downwards. Because we observe the level of trust *after* they have been cheated, this tends to generate a negative correlation between cheating and trust. When we run OLS estimates of (2) for the various domains we indeed find that this negative correlation is predominant (Appendix, Table A4).

To address this reverse causality issue, ideally we would need to observe the level of trust *before* people were cheated, which we do not. An alternative is to instrument current trust levels with variables that systematically affect an individual's propensity to trust others, but are unlikely to respond to shocks in being cheated. Since current trust depends on individuals' initial trust and on the experience of being cheated, if we can find variables that are correlated with initial trust we can identify the effect of trust on being cheated.

To obtain this exogenous source of variation we follow two strategies. First, in this subsection we use proxies for individual trustworthiness as instruments for trust as suggested by false consensus and supported by the experimental evidence provided in Section 8. Second, in the next subsection we use variation across immigrants in the average trust levels prevailing in their countries of origin.

Our instrument is a measure of how much responsibility is delegated to individuals by their supervisors at work. Specifically, the ESS asks individuals to state, on a scale from zero to 10, how much latitude their manager grants them along three different dimensions: a) freedom in organizing their daily work; b) power to influence policy decisions about the activities of the organization; and c) freedom to choose or change the pace of their work.<sup>25</sup>

We average the answers from the three parts of the delegation question to construct a single measure of how much authority individuals' managers grant them on the job. Since more trustworthy individuals are more likely, *ceteris paribus*, to be delegated more power and freedom of choice, we use this variable as a proxy for individuals' intrinsic trustworthiness. If individuals indeed extrapolate from their own type when forming trust beliefs, this index should have predictive power on measured trust. To be a valid instrument we also require that workplace delegation has no direct effect on individuals' risk of being cheated in the domains we observe. We see no obvious reason why such delegation would directly affect the chances that a person is cheated by, e.g., a mechanic or a plumber. Similarly, we do not

---

<sup>25</sup>The exact wording is: "Please say how much the management at your work allows/allowed you: a) ... to decide how your own daily work is/was organised; b) ... to influence policy decisions about the activities of the organisation; c) ... to choose or change your pace of work"

see how shocks to how frequently a person is cheated in his private life—which is private information and thus unobservable to the manager—could affect the amount of delegation a manager grants this person on the job. The only reason why there could be a correlation with the residuals in the cheating regression is because there could be an uncontrolled-for individual characteristic making it obvious to an outsider that the individual is susceptible to being cheated which would also reduce delegation to this individual. If this were the case, the IV estimates would be inconsistent. However, the inconsistency would take the form of a downwardly biased estimate of the true effect of trust on the frequency of being cheated. Since, as we will see, the IV estimates suggest a positive effect, this should be taken as a lower bound of the true effect of trust on the risk of being cheated.

Table 5 shows the first stage regression, focusing on the excluded instruments. Consistent with our identification strategy the instrument has a positive effect on the level of individual trust and is highly statistically significant.

Table 6 shows the results of the IV estimates. The first four columns report results for each of the four domains. In all cases, the negative effect of trust beliefs in the OLS estimates is reversed by the IV estimates, and a positive effect of trust on the number of times an individual has been cheated results. Economically, the effect of trust on exposure to cheating is substantial. Increasing trust by one standard deviation raises the number of times one is cheated by a bank by 1.5 (3 times the sample mean); the frequency of being cheated when buying second hand goods by 0.24 (62% of the sample mean); the frequency of being cheated when buying food by 1.3 (a bit more than the sample mean); and increases how frequently one is cheated by a plumber or repairer by 0.98 (1.7 times the sample mean). The remaining two columns show estimates using as the dependent variable the total number of times an individual was cheated on any domain (column 5) and the principal component of the measure of being cheated (column 6). In all cases the IV estimate shows a positive and highly significant effect of trust beliefs on being cheated.

The reduced form estimates of the effect of delegation (Table A5) imply that the effect of delegation on the number of times one is cheated is close to that implied by the first and the second stage of the IV estimates, lending indirect support to the validity of this instrument.

Overall, these estimates imply a large effect of trust on exposure to cheating. This is consistent with the idea that mistrust shields individuals from the risk of being cheated, while

too much trust amplifies this risk and hinders individual economic performance, lending support to one of the mechanisms through which heterogeneity in trust beliefs can produce a hump-shaped relationship between trust and income.

### 7.3 Immigrants, persistence and learning

The second strategy we use to identify the effect of heterogeneous trust beliefs on exposure to the social risk of being cheated exploits heterogeneity in values and beliefs among immigrants in a given country. The latter variation also allows us to provide evidence about how persistent (“wrong”) trust beliefs are.

We use information in the ESS about respondents’ country of origin as well as variation in average trust across countries. The recent literature on culture and economics shows that movers from one country to another tend to carry with them their cultural beliefs and norms (Giuliano (2007); Fernandez and Fogli (2009); Guiso et. al. (2006)) which they then transmit to their progeny in the new environment. Thus, either because priors instilled by parents are slow to change or because inherited norms and values, themselves, are slow to change, the cultural beliefs and norms prevailing in immigrants’ countries of origin continue to shape their beliefs. Consequently, people from high trust countries should be more exposed to the risk of being cheated than people from low trust countries. Hence we can use heterogeneity in average trust in immigrants’ countries of origin to identify the effect of beliefs on social risk. Since the level of trust in their country of origin is given, and cannot be affected by immigrants’ experiences of being cheated, we can exclude reverse causality due to learning. In effect, the average level of trust in an immigrant’s country of origin is a good measure of his or her initial trust prior.<sup>26</sup>

Of course, if there is learning the effect of the initial priors may fade away and disappear as the years in the new environment accumulate; or the effect of initial priors may be strong in the current generation, but fade in subsequent generations who grow up and learn in the new environment. Hence, by distinguishing between recent immigrants and immigrants

---

<sup>26</sup>The following story reported by Bhajan Grewal makes this point clearly. Several immigrants from Punjab, India worked in North Queensland, Australia, in sugarcane fields in the latter part of the 19th century. One of them—Mr Singh—was saving his earnings but keeping the money under the mattress. The Italian farmer took Mr Singh to Cairns, the largest regional town nearby and suggested to him to deposit his capital in a bank, which he did, reluctantly, because he had no experience with banks back home in India and because he did not trust the white fellas. After spending a few days with great unease, Singh went back to Cairns. When in town, he went into the bank, requested to withdraw his entire account balance, counted it carefully and deposited it all back after satisfying himself that the money was all there! The story illustrates very well the cultural influences on immigrants priors, especially on the first generation.

who arrived further in the past we can also shed light on how persistent trust mistakes (and false consensus) may be in reality.

To verify the persistence of the effects of initial trust beliefs, we restrict attention to the subsample of the ESS comprised only of first or second generation immigrants—leaving us with a sample of about 4,300 individuals. To each individual we associate the average level of trust in their country of origin (if first generation) or the average level of trust in their parents’ country of origin (if second generation). We then regress the number of times individuals have been cheated along our four domains on the average trust in the country of origin interacted with indicators for first- and second-generation immigrants and the usual controls, as well as dummies for first and second generation.

Table 7 shows the estimates. Trust in the country of origin has a positive, large and significant effect on the number of times a first-generation immigrant is cheated by a bank, by a plumber or repairer, or when buying food. Trust of origin is also positively related to how frequently individuals are cheated when purchasing second hand goods, but the effect is estimated with a large standard error.

When we use our aggregate measures of exposure to cheating (last two columns), the effect of trust in the country of origin for first-generation immigrants is always positive and significant. Overall, this evidence suggests that first-generation immigrants who move from high-trust countries are more likely to be victims of cheating than first-generation immigrants who move from low-trust countries. This is consistent with cultural priors exerting a life-long effect, perhaps because of false consensus.

In the second row we test whether the average level of trust in the country of ancestors’ origins still affects second-generation immigrants. We find no evidence that this is the case. For second-generation immigrants, average trust levels in the country of their ancestors’ origins have no effect on the number of times one is cheated in any of the specifications considered, nor in any cheating domains for which we have data.

The ESS also contains information on the number of years a person has been in the country, allowing us to further investigate whether the effect of trust on the frequency of being cheated attenuates as knowledge of the new environment is accumulated through experience. For this we focus on first-generation immigrants (about 2,000 observations) divided into two types: i) the newly-arrived (in the country for less than 20 years); and ii) “old arrivals” (in the country for at least 20 years). We interact the average trust level in

the country of origin with an indicator for an immigrant’s type. The estimates, reported in Table 8, show a large and significant effect of trust in the country of origin on the frequency with which newly-arrived immigrants are cheated by banks, by plumbers and by mechanics. The effect is still positive and sizable for the number of times the newly-arrived are cheated when buying goods second hand or when buying food, but the coefficients are not significant.<sup>27</sup> The effect is still positive across all domains for old arrivals, but typically smaller and never statistically significant.

All in all, this evidence shows that exogenous variation in beliefs about others’ trustworthiness affects individuals’ exposure to social risk. These effects can be quite persistent, as differences in priors that results from a tendency to extrapolate from received norms of trustworthiness can exert their effects for as long as 20 years.

## 8 Values, false consensus and heterogeneous trust beliefs

In this section we provide evidence of heterogeneity and persistence in trust beliefs, as well as a link to cultural norms through false consensus, in the context of a trust game (Berg, Dickhaut and McCabe (1995)) experiment. We also document the hump-shaped relationship between performance and beliefs in this more controlled context.

There are two reasons why it is important to document heterogeneity in beliefs in the context of a trust game experiment. First, it has been argued that trust questions of the sort asked in the ESS may measure not only individuals’ expectations about others’ trustworthiness, but also their preferences for risk and betrayal aversion (Fehr, 2009; Miller and Mitamura, 2003) or unconditional kindness (Ashraf et al., 2006). Hence, the heterogeneity in Figure 1 may be due to heterogeneity in preferences rather than heterogeneity in beliefs. For this reason, and in spite of our efforts to take preferences into account, one may be skeptical about the evidence in Section 5. In a trust game experiment we can elicit beliefs directly, allowing us to separate them from preferences. Second, one may suspect that the heterogeneity in trust shown in Figure 1 may be a reflection of heterogeneity in the pool

---

<sup>27</sup>One interpretation is that on goods that are traded more frequently and for which feedback on quality is faster—such as food or durables—learning is quicker and trust in the country of origin stops mattering earlier. On the other hand, learning may be slow for goods that are traded less frequently, such as repair services or financial services, and for which feedback on the quality of the service may be obtained infrequently and ambiguously. After all, investors in Madoff’s fund only discovered that they were cheated because of the financial crisis, and we get to know of a dishonest mechanic only after several trials, when chance can be ruled out.



of people respondents have in mind when answering the trust question. For instance, if respondents interpret the survey trust question as asking about the pool of people with whom they generally interact—who may well differ in their actual trustworthiness—then heterogeneity in responses to the trust question is not evidence about errors in trust beliefs. Everyone could simultaneously hold different, and correct, beliefs. Since our model is predicated on the effects of errors in trust beliefs, it would be helpful to document widespread heterogeneity in trust beliefs when individuals are asked to estimate the average trustworthiness of an unambiguously common population. We do this in the lab by eliciting participants’ beliefs about the trustworthiness of a well-defined set of people: other experimental participants.

## 8.1 The trust game experiment

For details of the experimental design, see Appendix 3. Here, we discuss the features necessary to understand the results.

Participants were undergraduates at LUISS Guido Carli University in Rome, Italy. Eight sessions were conducted, involving about 16 participants each, for a total sample of 124 participants.

Participants all played the trust game, which is a two-player sequential moves game of perfect information involving two roles: sender and receiver. Senders are endowed with 10.50 euros, while receivers have no endowment. Senders move first and decide whether to send some, all or none of their endowment to their receiver. Sending a positive amount entails a small fee of 0.50 euros and, after paying this fee, senders can send any positive whole-euro amount.<sup>28</sup> Each euro sent is tripled by the experimenter before being allocated to the receiver, who then decides whether to return some, all or none of this tripled amount to the sender.

We collected receivers’ actions using the strategy method: before discovering their co-players’ actions, receivers stated how much they would return for each amount they could possibly receive. Since there are ten possible amounts receivers can receive—3(=

---

<sup>28</sup>There are two slight departures here from the most-standard trust game (i.e., from the game in Berg, Dickhaut and McCabe, 1995). The first is the investment fee: the standard trust game does not involve an investment fee. The second is that our receivers receive no endowment, whereas in the standard trust game receivers’ endowments are equal to senders’ endowments. Neither departure changes the strategic nature of the game. For example, the unique subgame perfect equilibrium with standard (selfish) preferences still involves the sender sending nothing.

$1 \times 3), \dots, 30 (= 10 \times 3)$ —we collect a vector of ten actions for each receiver.<sup>29</sup> The amount the sender sends, together with the relevant element from the receiver’s strategy vector, determines the outcome of the trust game.

Since receivers could return any amount, including 0, returning a positive amount involves trustworthiness. We elicit all participants’ beliefs about this measure of trustworthiness in an incentive compatible manner and use them to construct a measure of trust beliefs independent from preferences.<sup>30</sup> Specifically, each participant estimated how much receivers would, on average, return for each amount receivers could receive, yielding 10 separate estimates per participant.<sup>31</sup> More accurate estimates were more highly remunerated, with perfect estimates paying 1 euro each. To construct a unidimensional measure of trust beliefs for each participant, we converted each of the 10 elements of his or her belief vector into percentage terms (0 to 1) and then took the average of these ten percentages. For example, suppose a participant’s belief vector is  $(1, 2, \dots, 10)$ —i.e., they believe that receivers will on average return 1 if they receive  $3 \times 1 = 3$ , 2 if they receive  $3 \times 2 = 6$ , etc. We would divide the first element by 3, the second by 6 and so on, to get the modified belief vector  $(\frac{1}{3}, \frac{2}{6}, \dots, \frac{10}{30})$  and then average the elements to get  $\frac{1}{3}$ , or 0.33, as the participant’s unidimensional trust belief.

To get a unidimensional measure of trustworthiness, for each receiver we apply the same procedures to their willingness-to-return vector. Consequently, we have a unidimensional trust belief measure for all participants, and a unidimensional trustworthiness measure for half of the participants for each round of game-play.

To study persistence of the relationship between trust beliefs and own trustworthiness, within each session the trust game was repeated for (up to) 12 rounds, with beliefs being elicited each round. To mitigate any nuisance dynamic effects such as reputation building, before each round subjects were anonymously and randomly re-paired and roles were randomly re-assigned within each pair. Furthermore, at the end of the experiment only one round of game-play was (randomly) chosen to count, with subjects being paid according to their decisions and beliefs in the chosen round only. This procedure eliminates the

---

<sup>29</sup>Precisely: how much a receiver will return if they receive 3 euros; ...; how much a receiver will return if they receive 30 euros.

<sup>30</sup>Individuals currently playing receiver were asked to exclude their own actions from the estimation and are paid for accuracy on this basis as well: i.e., when computing how accurate receivers’ beliefs are, each receiver’s own action is excluded from the calculation.

<sup>31</sup>Each participant was asked 10 questions: “how much, on average, will receivers return if they receive 3 euros?”; ...; “how much, on average, will receivers return if they receive 30 euros?”

possibility of wealth effects across rounds and is standard in the literature.

As a measure of “initial” trustworthiness mostly untainted by learning, we assign to each individual their unidimensional trustworthiness measure from the first time they played receiver, provided this occurred in one of the first two rounds.<sup>32</sup> Since roles are randomly re-assigned each round, this measure is defined for a large majority of participants, but not all of them (92 of 124).

Finally, all participants filled out a brief survey. The survey was sent (e-mailed) several days removed from laboratory sessions—a week before or after the participant’s session—to mitigate concerns that participants’ survey responses could systematically affect their decisions in the lab. As part of the survey, respondents were asked to report, on a scale from 0 to 10, how much emphasis their parents placed on a number of principles and behavioral rules during their upbringing (frugality, prudence, loyalty, etc.).<sup>33</sup> We use answers to (some of) these questions to construct a measure of the strength of received cultural values and norms of trustworthiness for each participant.

**Heterogeneity** Figure 7 shows the distribution of (unidimensional) trust beliefs in the first round of the trust game, when no learning about the trustworthiness of the pool of participants had yet been possible (panel A) and of our behavioral measure of own initial trustworthiness (panel B). Since trust beliefs and trustworthiness are measured by the average share that participants expect receivers will send back, and by the average share that receivers are willing to send back, respectively, these variables take values between 0 and 1. As these measures are continuous variables we report kernel density estimates. The figure documents considerable heterogeneity in trust priors, confirming the evidence in Figure 1. Since beliefs in the experiment refer to a common pool of people, heterogeneity in trust beliefs cannot be automatically ascribed to variation in the pools of people whose trustworthiness is being estimated.<sup>34</sup> Furthermore, since beliefs are measured independently of behavior, the heterogeneity in Figure 7, panel A, cannot reflect differences in risk

---

<sup>32</sup>The choice of the first two rounds balances concerns of contamination by learning which suggest only including those who were receivers in the first round—and leaving the measure undefined for half of the participants—against concerns about sample size which suggest extending the definition to include as many rounds as possible. In the end, we believe our definition is reasonable.

<sup>33</sup>A wide array of questions was asked, some completely irrelevant to trust and trustworthiness, in order to mitigate experimenter/demand effect in the survey answers and in the experiment.

<sup>34</sup>It is true that the Figure 7, panel A, reports beliefs for all sessions pooled, so some people might still question the source of heterogeneity. However, plotting the trust belief densities for each session separately (not reported, but available upon request) also yields quite a lot of heterogeneity.

attitudes.<sup>35</sup> In the sample the average level of trust beliefs is 0.27.<sup>36</sup> The standard deviation of trust beliefs is 0.16.

The figure also documents substantial heterogeneity in behavioral trustworthiness, whose sample mean and standard deviation are 0.32 and 0.16, respectively. In the next section we test whether heterogeneity in trustworthiness is reflected in heterogeneous beliefs.

**False consensus and persistence** Table 9, panel A, shows regressions of trust beliefs in various rounds on own initial trustworthiness. To isolate, as best as possible, trustworthiness as an individual trait, we use initial trustworthiness as a regressor. To reduce sampling variation due to small sample size we aggregate observations over blocks of three rounds. As the first column shows, in early rounds initial trustworthiness is strongly positively correlated with trust beliefs, lending support to the idea that individuals form beliefs about others' trustworthiness by extrapolating from their own types. Quite remarkably, own trustworthiness explains about 60% of the initial heterogeneity in beliefs. As the second column shows, this tendency does not vanish when the game is repeated and people are thus given the opportunity to learn about the pool of participants. The correlation weakens, and the effect is somewhat smaller, in later rounds but both remain sizable and significant. Thus, initial trustworthiness still affects trust beliefs even after the game has been played several times, always drawing from an invariant pool of individuals, which we take as evidence that false consensus persists. However, the decline in the strength of the link also suggests that given enough opportunities to learn about a stable pool of people, the tendency to attribute to others one's own trustworthiness may vanish.<sup>37</sup>

This evidence is consistent with the idea that priors are driven, through false consensus, by norms of behavior that shape individual's own trustworthiness. To make this link even more clear and show the ultimate relationship between cultural values and beliefs we use information on the moral values emphasized by participants' parents. For our purposes, we use parents' emphasis on two values: the first is how much emphasis an individual's

---

<sup>35</sup>Unless the elicitation procedure is biased by risk preferences as well. We cannot rule this out completely, as how to do so is a still-unsettled debate within experimental economics. We use a very standard quadratic scoring rule. There is experimental evidence suggesting that this mechanism elicits beliefs reasonably accurately regardless of risk preferences (see, e.g., Huck and Weiszäcker, 2002).

<sup>36</sup>Since every dollar sent is tripled, 0.33 would imply senders believe that receivers will return as much as is sent.

<sup>37</sup>An interesting question is whether the false consensus effect reappears any time an individual faces a new pool of people or the pool she is interacting with changes.

parents placed on teaching to always behave as good citizens; the second is the emphasis parents placed on loyalty to groups or organizations. We average the responses to these two questions and divide the result by 10 to put it on a scale—0 to 1—comparable with beliefs. We use this measure as a proxy for individuals’ intrinsic trustworthiness, an individual-specific trait.

Table 9, Panel B shows that this measure of parents’ effort spent on teaching good values is correlated with individuals’ initial trustworthiness, which is consistent with behavioral types reflecting heterogeneous cultural values. Of course, it is imperfectly correlated, partly because the measure of values that we have is only a proxy for the true trait, and partly because own traits are also shaped by interactions in the social sphere. Panel C shows direct regressions of trust beliefs on our survey measure of cultural values: at all repetitions the cultural measure of trustworthiness predicts trust beliefs.

In sum, the evidence in this section shows three things. First, when no information is available about a group, individuals form beliefs about the trustworthiness of others extrapolating from their own types, which are quite heterogeneous. Second, this tendency is highly persistent, though attenuated through learning. Third, heterogeneity in own trustworthiness can be traced back to heterogeneous cultural norms instilled by parents. Measures of the latter can provide valuable instruments for trust beliefs, an implication that justifies our choice of instruments in Section 7.

**Beliefs and performance** To study the relationship between trust beliefs and economic performance we rely on a modified version of the trust game. In particular, to match our model (Section 3), each dollar sent in this modified trust game is increased according to a concave “production” function before reaching the receiver.<sup>38</sup> Concavity guarantees an internal optimal send amount given a sufficiently large proportion of trustworthy receivers. We found this last condition was only satisfied in the first round of a pilot study using multiple rounds of this concave trust game. Therefore, the experiment reported in this section uses a single-round design.

There were two additional modifications. First, we conducted the experiment on-line to minimize experimenter/demand effect.<sup>39</sup> Second, we used a full strategy method: par-

---

<sup>38</sup>If the sender sends  $S$  euros, the receiver receives  $8S^{0.5}$  euros.

<sup>39</sup>Conducting a single-round experiment on-line is much easier than conducting a multiple-round experiment on-line as it requires much less real-time feedback. This is one reason we chose not to implement the

ticipants submitted their actions for each role before knowing which role they would be assigned. This allows us to construct a sender performance measure for all participants since we know what each participant would do as a sender. For more details on the experimental design, see Appendix 3.

The on-line trust game was played by 122 students at LUISS over 4 separate days.<sup>40</sup> Each day constitutes a session. For each participant,  $i$ , we construct a measure of performance by randomly choosing another participant,  $j$ , from the same experimental session and computing  $i$ 's earnings using  $i$ 's sender strategy and  $j$ 's receiver strategy.<sup>41</sup>

We use willingness-to-return amounts—excluding each participant's own actions—and beliefs about these return amounts within each session to construct a unidimensional measure of belief errors for each participant. Specifically, for each participant we first compute a separate belief error in percentage terms for each amount a receiver could have received. This yields ten separate belief error measures for each participant, each ranging from  $-1$  to  $1$ , where negative values indicate under-estimating. We use the average of these ten measures for each participant as our (unidimensional) belief errors measure, which again ranges from  $-1$  to  $1$ .

Figure 8 presents a scatter plot of the relationship between this belief errors measure and our performance measure in the on-line experiment. Consistent with our model and with the survey data, earnings are hump-shaped in belief errors. Both those who hold overly pessimistic trust beliefs (negative belief errors) and those who hold overly-optimistic trust beliefs (positive belief errors) earn less than those whose belief errors are approximately zero. This humped shape is confirmed by the regression presented in Table 10: the coefficient on the squared belief errors is both negative and significant.<sup>42</sup> Furthermore, the coefficient on the linear term, regardless of significance, implies that performance attains its maximum for belief errors close to zero. The estimated relationships suggest that senders earn between 11.00 and 11.45 euros on average when belief errors are zero, constituting a 5 to 9 percent

---

multiple-round experiment of the previous section on-line.

<sup>40</sup>There was no overlap in participants with the laboratory experiment.

<sup>41</sup>That is, performance for participant  $i$  is measured as the earnings they would have made if they had been assigned the role of sender:  $Y_i = 10.5 - S_i + \gamma_j 8S_i^{0.5} - 0.5I(S_i)$ , where  $\gamma_j$  denotes the proportion of the amount received,  $8S_i^{0.5}$ , what the receiver  $j$  paired with  $i$  returns and  $I(S_i)$  is an indicator function equal to 1 if  $i$  sends a positive amount.

<sup>42</sup>This continues to be true when we add dummies for each session to control for session fixed effects and when standard errors are clustered by session, where each separate day the experiment was conducted constitutes a session.

increase over the safe return (10.50 euros) from sending nothing.

To get another measure of the magnitude of income differences implied by belief errors we divided the data into three categories: “under-estimators,” “over-estimators,” and “accurate-estimators.” Accurate-estimators had belief errors within a small interval around zero,  $[-0.1, 0.1]$ ; under-estimators had belief errors below this interval; over-estimators had belief errors above this interval. Table 11 shows that accurate-estimators earned about 18 percent more on average than under-estimators, who, in turn, earned about the same as over-estimators.<sup>43</sup>

Summing up, the on-line experiment allows us to investigate the relationship our model predicts between economic performance and errors in trust beliefs. The results are consistent with both our model and the ESS evidence. Moreover, in terms of percent of income foregone, earnings differences in the on-line experiment are of the same order of magnitude as in the survey data.

## 9 Conclusions

We document the existence of a hump-shaped relationship between individual trust and individual income. For an individual the cost of miscalibrated trust beliefs can be substantial and of the same order of magnitude as returns to education. Though both excessive trust and excessive mistrust are individually costly, the data suggest that the cost of trusting too little far exceeds that of trusting too much, even in low trust countries. From a societal point of view, however, there is an important difference between the two excesses. While excessive mistrust and excessive trust are both *individually* costly, mistrust is also *socially* costly as it reduces surplus creation. On the contrary, excessive trust may create social surplus even if this surplus is allocated in a way that harms the overly trusting individual. This difference reconciles our findings of a concave relationship between performance and trust at the individual level and the monotonically increasing relationship found in aggregate data.

---

<sup>43</sup> As rough robustness checks (not reported, but available on request) we also ran the regressions in Table 11 using a wider interval— $[-0.15, 0.15]$ —or a narrower interval— $[-0.05, 0.05]$ —to define accurate-estimators, as well as using a definition of over- and under-estimators defined by the 33rd and 66th percentiles of the observed belief errors. None of these modifications changed the results qualitatively: accurate-estimators consistently earned more, on average, than others.

## References

- [1] Algan, Yann and Pierre Cahuc (2010) "Inherited Trust and Growth," *American Economic Review*, forthcoming.
- [2] Alesina, Alberto and Eliana La Ferrara (2002), "Who Trusts Others?," *Journal of Public Economics*, 85(2), 207-234.
- [3] Arrow, Kenneth (1972), "Gifts and Exchanges," *Philosophy & Public Affairs*, 1(4), 343-362.
- [4] Ashraf, Nava, Iris Bohnet and Nikita Piankov (2006), "Decomposing Trust and Trustworthiness," *Experimental Economics*, 9 (3), 193-208.
- [5] Berg, J., Dickhaut, J. and K. McCabe (1995), "Trust, Reciprocity and Social History," *Games and Economic Behavior*, 10, 122-142.
- [6] Bisin, Alberto, Giorgio Topa, and Thierry Verdier (2004). "Cooperation as a Transmitted Cultural Trait," *Rationality and Society*, 16 (4), 477-507.
- [7] Bisin, Alberto, and Thierry Verdier (2000). "Beyond the Melting Pot: Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits," *Quarterly Journal of Economics*, 115 (3), 955-988.
- [8] Blanco, Mariana (2008), "Within-Subject Experiments on Other-Regarding Preferences," Chapter 1, Phd thesis, Royal Holloway, University of London.
- [9] Bloom, Nicholas, Raffaella Sadun and John Van Reenen (2009), "The Organization of Firms across Countries," CEPR DP 937.
- [10] Camerer, Colin F (2003). *Behavioral Game Theory: experiments in strategic interaction*. Princeton, NJ: Princeton University Press.
- [11] Cox, James C. (2004), "How to Identify Trust and Reciprocity," *Games and Economic Behavior*, 46, 260-281
- [12] Dixit, Avinash (2003), "Trade Expansion and Contract Enforcement," *Journal of Political Economy*, 111 (6), 1293-1317.



- [13] Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde (2007). “The Intergenerational Transmission of Risk and Trust Attitudes,” IZA DP 2380.
- [14] Fehr, Ernst (2009), “On the Economics and Biology of Trust,” Presidential Address, European Economic Association, *Journal of the European Economic Association*, forthcoming
- [15] Fernández, Raquel and Alessandra Fogli (2009), “Culture: An Empirical Investigation of Beliefs, Work, and Fertility,” *American Economic Journal: Macroeconomics*, 1(1), 146-177.
- [16] Giuliano, Paola (2007), “Living Arrangements in Western Europe: Does Cultural Origin Matter?,” *Journal of the European Economic Association*, 5 (5), 927-952.
- [17] Glaeser, Edward, David Laibson, José A. Scheinkman and Christine L. Soutter (2000), “Measuring Trust,” *Quarterly Journal of Economics* 115(3), 811-846.
- [18] Guiso, Luigi, Paola Sapienza and Luigi Zingales (2004), “The Role of Social Capital in Financial Development,” *The American Economic Review*, 94, 526-556.
- [19] Guiso, Luigi, Paola Sapienza and Luigi Zingales (2006), “Does Culture Affect Economic Outcomes?,” *Journal of Economic Perspectives*, 20 (2), 23-48.
- [20] Guiso, Luigi, Paola Sapienza and Luigi Zingales (2008a), “Trusting the Stock Market,” *Journal of Finance*, 63 (6), 2557-2600.
- [21] Guiso, Luigi, Paola Sapienza and Luigi Zingales (2008b), “Social Capital as Good Culture,” *Journal of the European Economic Association*, 6(2-3), 295-320.
- [22] Guiso, Luigi, Paola Sapienza and Luigi Zingales (2009), "Cultural Biases in Economic Exchange?," *Quarterly Journal of Economics*, 124 (3), 1095-1131.
- [23] Huck, Steffen and Weizsäcker, Georg (2002), “Do Players Correctly Estimate What Others Do?: Evidence of Conservatism in Beliefs,” *Journal of Economic Behavior & Organization*, 47 (1), 71-85.
- [24] Knack, Stephen and Keefer, Philip (1996) “Does Social Capital Have an Economic Payoff? A Cross-Country Investigation,” *The Quarterly Journal of Economics*, 112 (4), 1251-1288.

- [25] Knack, Stephen and Paul Zak (2001) "Trust and Growth," *The Economic Journal*, 111, 295-321.
- [26] Krueger, Joachim and Russel W. Clement (1994), "The Truly False Consensus Effect: An Ineradicable and Egocentric Bias in Social Perception", *Journal of Personality and Social Psychology of Addictive Behaviors*, 67 (4):596-610.
- [27] La Porta, Rafael, Florencio Lopez de Silanes, Andrea Shleifer and Robert Vishny (1997) "Trust in Large Organizations," *The American Economic Review*, 87 (2), 333-338.
- [28] Luttmer, Erzo and Monica Singhal (2010), "Culture, Context, and the Taste for Redistribution", *American Economic Journal: Economic Policy*, forthcoming.
- [29] Massey, Cade and Richard H. Thaler (2006), "The Loser's Curse: Overconfidence vs. Market Efficiency in the National Football League Draft," University of Chicago, mimeo.
- [30] Miller, A. S., and T. Mitamura (2003). "Are Surveys on Trust Trustworthy?," *Social Psychology Quarterly*, 66, 62-70.
- [31] Nunn, Nathan and Leonard Wantchekon (2009), "The Slave Trade and the Origins of Mistrust in Africa," Harvard mimeo.
- [32] Ross, Lee, Greene, D., and House, P. (1977), "The False Consensus Phenomenon: An Attributional Bias in Self-Perception and Social Perception Processes," *Journal of Experimental Social Psychology*, 13(3), 279-301.
- [33] Sapienza, Paola, Anna Toldra and Luigi Zingales (2007), "Understanding Trust," NBER WP 13387
- [34] Schelling, T. (1966). *Arms and Influence*. New Haven, CT: Yale University Press.
- [35] Tabellini, Guido (2008a). "The Scope of Cooperation: Values and Incentives," *Quarterly Journal of Economics*, 123 (3), 905-950.
- [36] Tabellini, Guido (2008b), "Institutions and Culture," *Journal of the European Economic Association*, 6(2-3), 255-294.

- [37] Westen, Drew, Pavel S. Blagov, Keith Harenski, Clint Kilts and Stephan Hamann (2006), “Neural Bases of Motivated Reasoning: An MRI Study of Emotional Constraints on Partisan Political Judgment in the 2004 U.S. Presidential Election,” *Journal of Cognitive Neuroscience*, 18(11), 1947–58.

**Table 1**  
**Descriptive statistics**

A. European Social Survey

Variable	Whole sample		Immigrants sample	
	Mean	Std. Dev.	Mean	Std. Dev.
Log income	9.695	1.072	9.659	1.233
Trust	5.033	2.498	4.988	2.479
Risk tolerance	2.982	1.429	2.991	1.454
Age	46.691	17.771	46.044	17.255
Male	0.473	0.499	0.457	0.498
Immigrant	0.075	0.263		
Married	0.558	0.497	0.542	0.498
Single	0.263	0.440	0.256	0.436
Father primary	0.406	0.491	.370	0.483
Father secondary	0.487	0.500	0.474	0.499
Primary	0.192	0.394	0.163	0.369
Secondary	0.605	0.489	0.608	0.488
Big city	0.323	0.468	0.378	0.485
Small city	0.308	0.462	0.303	0.459
Unemployed	0.051	0.220	0.055	0.229
Out of labor force	0.535	0.499	0.531	0.499
Altruism	5.052	0.886		
Religiosity	4.97	2.98		
Cheated: Bank	1.494	0.892	1.466	0.869
Cheated: Second hand goods	1.374	0.792	1.357	0.772
Cheated: Food	2.106	1.435	2.085	1.384
Cheated: Plumber, repairer	1.629	1.000	1.685	1.043
Cheated (sum)	2.473	2.753	2.552	2.721
Cheated (principal component)	0.099	1.379	0.133	1.366
Trustworthiness	16.177	9.454		
Professionals	0.137	0.344		
Technicians	0.164	0.370		

Clerks	0.109	0.312		
Workers	0.142	0.349		
Agricultural workers	0.041	0.197		
Mechanics, repairers, textile workers	0.137	0.344		
Assemblers, operators and drivers	0.079	0.270		
Labourers and elementary occupations	0.103	0.303		
Fraction of first generation immigrants			0.535	0.499
First generation immigrants: old arrivals			0.534	0.499
Trust in the country of origin			0.271	0.132

#### B. Laboratory Experiment

Variable	mean	St dev
Good Values	0.637	0.199
Initial own trustworthiness	0.32	0.162
Expected trustworthiness (trust belief)	0.265	0.158
Return Proportion	0.211	0.18
Invest Amount	5.258	3.107
Invest Propensity	0.676	0.469

#### C. On-line Experiment

Variable	mean	St dev
Invest Propensity	0.730	0.446
Invest Amount	3.934	3.315
Estimates of Return Proportion	1.287	0.578
Return Proportion	1.312	0.669
Trust Belief Error	-0.007	0.145
Sender Earnings	10.950	3.077

**Table 2**  
**Trust and Income**

	(1)	(2)	(3)	(4)
	Log inc.	Log inc.	Log inc.	Log inc.
Trust1	0.003 (0.026)	0.004 (0.027)	0.006 (0.028)	
Trust 2	0.031 (0.023)	0.039* (0.024)	0.035 (0.024)	
Trust 3	0.071*** (0.021)	0.081*** (0.022)	0.086*** (0.022)	
Trust 4	0.082*** (0.022)	0.083*** (0.023)	0.081*** (0.023)	
Trust 5	0.081*** (0.020)	0.083*** (0.020)	0.085*** (0.021)	
Trust 6	0.119*** (0.022)	0.126*** (0.022)	0.124*** (0.023)	
Trust 7	0.134*** (0.021)	0.142*** (0.021)	0.142*** (0.022)	
Trust 8	0.138*** (0.022)	0.145*** (0.022)	0.145*** (0.022)	
Trust 9	0.133*** (0.027)	0.138*** (0.028)	0.141*** (0.028)	
Trust 10	0.071** (0.035)	0.079** (0.036)	0.091** (0.037)	
Trust				0.030*** (0.006)
Trust squared				-0.002** (0.001)
Age	0.002* (0.001)	0.003** (0.002)	0.003** (0.002)	0.003** (0.002)
Age squared	-0.000*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)
Male	0.049*** (0.008)	0.041*** (0.008)	0.040*** (0.009)	0.041*** (0.008)
Primary education	-0.483*** (0.015)	-0.480*** (0.015)	-0.478*** (0.016)	-0.482*** (0.015)
Secondary education	-0.262*** (0.010)	-0.261*** (0.011)	-0.262*** (0.011)	-0.263*** (0.011)

Married	0.408*** (0.011)	0.407*** (0.012)	0.408*** (0.012)	0.407*** (0.012)
Single	0.115*** (0.015)	0.118*** (0.016)	0.121*** (0.016)	0.118*** (0.016)
Father primary education	-0.155*** (0.016)	-0.157*** (0.016)	-0.155*** (0.017)	-0.157*** (0.016)
Father secondary education	-0.054*** (0.014)	-0.057*** (0.015)	-0.057*** (0.015)	-0.057*** (0.015)
Unemployed	-0.508*** (0.021)	-0.500*** (0.021)	-0.506*** (0.022)	-0.501*** (0.021)
Out of labor force	-0.179*** (0.010)	-0.167*** (0.010)	-0.166*** (0.010)	-0.167*** (0.010)
Immigrant	-0.158*** (0.016)	-0.157*** (0.016)	-0.158*** (0.016)	-0.158*** (0.016)
Big city	0.065*** (0.011)	0.072*** (0.011)	0.073*** (0.011)	0.073*** (0.011)
Small city	0.031*** (0.010)	0.039*** (0.010)	0.037*** (0.010)	0.040*** (0.010)
Risk tolerance		0.015*** (0.003)	0.014*** (0.003)	0.015*** (0.003)
Altruism			-0.019*** (0.005)	
Observations	30254	28383	27687	28383
R-squared	0.62	0.62	0.62	0.62
Test trust2= trust8 (p-values)	0.00	0.00	0.00	
Test trust10=trust8 (p-values)	0.05	0.04	0.03	

Notes: [1] Each regression controls for country and region fixed effects; the F-test for the joint significant of these coefficients has always p-value of 0.000; [2] Standard errors are clustered at the country level, \*\*\* significant at 1%, \*\* significant at 5%, \* significant at 10%. [3] *Trust* is the answer to the following question: “Generally speaking would you say that most people can be trusted or that you can’t be too careful in dealing with people? Please tell me on a score of 0 to 10, where 0 means you can’t be too careful and 10 means that most people can be trusted”; *Risk tolerance* and *Altruism* are the answers to the following questions: “Now I will briefly describe some people. Tell me how much each person is or is not like you: very much like me (6), like me (5), somewhat like me (4), a little like me (3), not like me (2), not like me at all (1); She/he looks for adventures and likes to take risks. She/he wants to have an exciting life. (*Risk aversion*); It is important to her/him to be loyal to her/his friends. She/he wants to devote herself/himself to people close to her/him. (*Altruism*). [4] The excluded group for education and father education are people with college or more; the excluded group for marital status is divorced or widows/widowers; the excluded group for labor status are people employed; the excluded group for city size are people living in a country village, a farm or a home in the countryside.

**Table 3**  
**Trust and Income, Controlling for Moderation**

	1	2	3	4	5	6	7	8	9	10
Trust	0.003 (0.027)	0.037 (0.024)	0.079*** (0.022)	0.081*** (0.023)	0.081*** (0.020)	0.124*** (0.022)	0.139*** (0.021)	0.143*** (0.022)	0.136*** (0.028)	0.079** (0.036)
Risk tolerance	0.025** (0.012)	0.047*** (0.014)	0.066*** (0.015)	0.069*** (0.017)	0.054** (0.023)					
Test of equality of coef. (p-values)	All equal coeff. (0.000)		Trust2=Trust8 (0.000)		Trust10=Trust8 (0.05)					
Trust	0.003 (0.027)	0.036 (0.024)	0.079*** (0.022)	0.079*** (0.023)	0.080*** (0.020)	0.124*** (0.022)	0.139*** (0.021)	0.143*** (0.022)	0.135*** (0.028)	0.076** (0.036)
Altruism	0.086 (0.083)	0.095 (0.076)	0.089 (0.074)	0.110 (0.074)	0.098 (0.074)					
Test of equality of coef. (p-values)	All equal coeff. (0.000)		Trust2=Trust8 (0.000)		Trust10=Trust8 (0.04)					
Trust	-0.011 (0.030)	0.022 (0.026)	0.060** (0.024)	0.058** (0.025)	0.065** (0.023)	0.103*** (0.024)	0.114*** (0.023)	0.123*** (0.024)	0.112*** (0.030)	0.058 (0.039)
Political ideology	-0.030 (0.039)	0.035 (0.031)	-0.008 (0.029)	0.001 (0.029)	0.010 (0.027)	0.044 (0.029)	0.084*** (0.029)	0.074** (0.030)	0.056 (0.036)	0.024 (0.036)
Test of equality of coef. (p-values)	All equal coeff (0.000)		Trust2=Trust8 (0.000)		Trust10=Trust8 (0.06)					
Trust	0.002 (0.027)	0.033 (0.024)	0.075*** (0.022)	0.077*** (0.023)	0.079*** (0.020)	0.119*** (0.022)	0.137*** (0.021)	0.143*** (0.022)	0.142*** (0.028)	0.084** (0.036)
Religiosity	0.030 (0.021)	0.005 (0.019)	0.048*** (0.018)	-0.007 (0.020)	-0.023 (0.016)	-0.014 (0.018)	-0.017 (0.017)	-0.053*** (0.018)	-0.077*** (0.022)	-0.095*** (0.022)
Test of equality of coef. (p-values)	All equal coeff (0.000)		Trust2-Trust8 (0.07)		Trust10=Trust8 (0.07)					

Notes: [1] Each regression controls for country, region fixed effects and the full set of controls of column 3 Table 2; [2] Standard errors are clustered at the country level, \*\*\* significant at 1%, \*\* significant at 5%, \* significant at 10%. [3] *Trust* is the answer to the following question: “Generally speaking would you say that most people can be trusted or that you can’t be too careful in dealing with people? Please tell me on a score of 0 to 10, where 0 means you can’t be too careful and 10 means that most people can be trusted”; *Risk tolerance* and *Altruism* are the answers to the following question: “Now I will briefly describe some people. Tell me how much each person is or is not like you: very much like me (6), like me (5), somewhat like me (4), a little like me (3), not like me (2), not like me at all (1); She/he looks for adventures and likes to take risks. She/he wants to have an exciting life. (*Risk aversion*); It is important to her/him to be loyal to her/his friends. She/he wants to devote herself/himself to people close to her/him. (*Altruism*); *Religiosity* is a question asking the respondents how religious they are; the scale goes from “Not at all religious” (0) to “Very religious” (10). *Political ideology* is a question asking the respondents their political ideology; the scale goes from “Left” (0) to “Right” (10).



**Table 4**  
**Trust and Income in Low and High Trust Countries**

	Log income Trust lower than median	Log income Trust above or equal to median	Log income Low trust	Log income High trust
Trust 1	0.011 (0.034)	-0.001 (0.044)	0.008 (0.033)	0.007 (0.047)
Trust 2	-0.005 (0.031)	0.105*** (0.037)	0.006 (0.029)	0.103*** (0.039)
Trust 3	0.069** (0.029)	0.100*** (0.033)	0.060** (0.028)	0.113*** (0.035)
Trust 4	0.058* (0.032)	0.104*** (0.034)	0.058* (0.030)	0.109*** (0.036)
Trust 5	0.085*** (0.028)	0.086*** (0.031)	0.076*** (0.027)	0.095*** (0.033)
Trust 6	0.140*** (0.034)	0.130*** (0.032)	0.143*** (0.032)	0.130*** (0.034)
Trust 7	0.144*** (0.034)	0.154*** (0.031)	0.150*** (0.032)	0.156*** (0.033)
Trust 8	0.116*** (0.040)	0.162*** (0.032)	0.120*** (0.038)	0.165*** (0.034)
Trust 9	0.053 (0.063)	0.163*** (0.037)	0.049 (0.060)	0.167*** (0.038)
Trust 10	0.066 (0.078)	0.093** (0.044)	0.063 (0.074)	0.097** (0.046)
Observations	9971	18412	10916	17467
R-squared	0.42	0.51	0.46	0.51
Test of equality of coeff. (p-values)	tru7=tru9 (0.14) tru7=tru2 (0.00)	tru9=tru10 (0.07) tru9=tru2 (0.07)	tru7=tru9 (0.09) tru7=tru2 (0.00)	tru9=tru10 (0.08) tru9=tru2 (0.05)

Notes: [1] The specification follows the main specification of Table 2; [2] Standard errors are clustered at the country level, \*\*\* significant at 1%, \*\* significant at 5%, \* significant at 10%. [3] *Trust* is the answer to the following question: “Generally speaking would you say that most people can be trusted or that you can’t be too careful in dealing with people? Please tell me on a score of 0 to 10, where 0 means you can’t be too careful and 10 means that most people can be trusted”; [4] Column 1 restrict the sample to countries with trust lower than the median in the sample (equal to 5); column 2 restricts the sample to countries with trust equal or higher than the median; [5] In columns 3, we define low trust countries those with fat tail to the left including Greece, Italy, Portugal, Turkey, Czech Republic, Hungary, Poland, Slovenia and Slovakia; in column 4, we define high trust countries those with fat tails to the right, including Switzerland, Denmark, Finland, Iceland, Netherland, Sweden and Norway.

**Table 5**  
**Trust and Cheating: First Stage Regressions**

	(1): Bank insurance	(2) Second hand things	(3) Food	(4) Plumber, builder, mechanic, repairer	(5) Times being cheated (sum)	(6) Being cheated (principal component)
Trustworthiness	.0084*** (.0020)	.0078*** (.0019)	.0082*** (.0019)	.0087*** (.0019)	.0089*** (.0021)	.0089*** (.0021)
Observations	21163	22663	23062	22463	19774	19774
R-squared	0.23	0.23	0.23	0.23	0.24	0.23

Notes: [1] Standard errors are clustered at the country level, \*\*\* significant at 1%, \*\* significant at 5%, \* significant at 10%. [2] *Trust* is the answer to the following question: “Generally speaking would you say that most people can be trusted or that you can’t be too careful in dealing with people? Please tell me on a score of 0 to 10, where 0 means you can’t be too careful and 10 means that most people can be trusted”; *Trustworthiness* is the sum of the following three questions: “I am going to read out a list of things about your working life. Using this card, please say how much the management at your work allows/allowed you to 1) decide how your own daily work is/was organized; 2) influence policy decisions about the activities of the organization? 3) to choose or change your pace of work?” The answer to each question can take values from I have/had no influence (0) to I have had complete control (10). [3] The cheating variables are the answer to the following questions: “how often, if ever, have each of these things happened to you in the last five years? A bank or insurance company failed to offer you the best deal you were entitled to; you were sold something second-hand that quickly proved to be faulty; you were sold food that was packed to conceal the worse bits; A plumber, builder, car mechanic or other repair person overcharged you or did unnecessary work” The answer could take values Never (1), once (2), twice (3), 3 or 4 time (4), 5 times or more (5).

**Table 6**  
**Trust and Cheating, Instrumental Variable Regressions**

	(1) Bank insurance	(2) Second hand things	(3) Food	(4) Plumber, builder, mechanic, repairer	(5) Times being cheated (sum)	(6) Being cheated (principal component)
Trust	0.817*** (0.219)	0.234** (0.100)	0.599*** (0.197)	0.534*** (0.158)	2.271*** (0.610)	1.164*** (0.312)
Age	0.010* (0.006)	-0.008*** (0.003)	0.013*** (0.005)	0.008* (0.004)	0.016 (0.017)	0.007 (0.009)
Age squared	-0.000*** (0.000)	-0.000 (0.000)	-0.000*** (0.000)	-0.000*** (0.000)	-0.001*** (0.000)	-0.000*** (0.000)
Male	0.099*** (0.033)	0.088*** (0.015)	-0.173*** (0.029)	0.112*** (0.024)	0.128 (0.096)	0.106** (0.049)
Immigrant	0.009 (0.057)	0.043* (0.026)	-0.004 (0.051)	0.033 (0.044)	0.046 (0.171)	0.028 (0.087)
Married	-0.160*** (0.059)	-0.059** (0.026)	-0.108** (0.054)	-0.174*** (0.045)	-0.538*** (0.178)	-0.279*** (0.091)
Single	-0.279*** (0.069)	-0.047 (0.031)	-0.235*** (0.064)	-0.254*** (0.052)	-0.795*** (0.199)	-0.403*** (0.102)
Primary	0.214** (0.109)	0.103** (0.048)	0.114 (0.099)	0.117 (0.079)	0.662** (0.319)	0.345** (0.163)
Secondary	0.202** (0.081)	0.099*** (0.037)	0.090 (0.072)	0.090 (0.058)	0.573** (0.238)	0.304** (0.122)
Father primary	0.236** (0.100)	0.023 (0.045)	0.110 (0.090)	0.054 (0.072)	0.379 (0.267)	0.199 (0.136)
Father secondary	0.155** (0.076)	0.017 (0.035)	0.030 (0.066)	0.045 (0.057)	0.267 (0.210)	0.140 (0.107)
Unemployed	0.250*** (0.092)	0.112*** (0.043)	0.249*** (0.085)	0.161** (0.067)	0.692*** (0.258)	0.350*** (0.132)
Out of labor force	0.056 (0.036)	0.035** (0.017)	0.167*** (0.035)	0.071** (0.029)	0.283*** (0.105)	0.131** (0.053)
Risk tolerance	-0.009 (0.015)	0.006 (0.007)	-0.031** (0.014)	-0.001 (0.011)	-0.043 (0.044)	-0.018 (0.023)
Log income	-0.049 (0.033)	-0.036** (0.015)	-0.030 (0.031)	-0.017 (0.025)	-0.133 (0.092)	-0.071 (0.047)
Big city	0.095** (0.044)	0.024 (0.020)	0.195*** (0.039)	0.113*** (0.033)	0.481*** (0.125)	0.230*** (0.064)
Small city	0.123*** (0.046)	0.058*** (0.020)	0.166*** (0.039)	0.114*** (0.032)	0.489*** (0.130)	0.244*** (0.067)
Observations	21163	22633	23062	22463	19774	19774

Notes: [1] Each regression controls for country, region fixed effects and 8 occupational dummies [2] Standard errors are clustered at the country level, \*\*\* significant at 1%, \*\* significant at 5%, \* significant at 10%. [3] *Trust* is the answer to the following question: “Generally speaking would you say that most people can be trusted or that you can’t be too careful in dealing with people? Please tell me on a score of 0 to 10, where 0 means you can’t be too careful and 10 means that most people can be trusted”; [4] *Trust* is instrumented using the variable *Trustworthiness* obtained as the sum of the answers to the following three questions: “I am going to read out a list of things about your working life. Using this card, please say how much the management at your work allows/allowed you to 1) decide how your own daily work is/was organized; 2) influence policy decisions about the activities of the organization? 3) to choose or change your pace of work?” The answer to each question can take values from 0 (I have/had no influence) to 10 (I have/ had complete control). [5] The excluded group for education and father education are people with college or more; the excluded group for marital status is “divorced or widow/er”; the excluded group for labor status are people employed; the excluded group for city size are people living in a country village, a farm or a home in the countryside.

**Table 7**  
**Trust and Cheating, Evidence from First and Second Generation Immigrants**

	(1)	(2)	(3)	(4)	(5)	(6)
	Bank insurance	Second hand things	Food	Plumber, builder, mechanic, repairer	Times being cheated (sum)	Being cheated (principal component)
Trust-first gen.	0.271** (0.103)	0.080 (0.154)	0.666*** (0.220)	0.348* (0.195)	1.491*** (0.489)	0.698*** (0.240)
Trust-second gen.	-0.031 (0.194)	0.127 (0.211)	-0.224 (0.171)	-0.066 (0.265)	-0.493 (0.614)	-0.218 (0.326)
Risk tolerance	0.031*** (0.011)	0.030** (0.014)	0.004 (0.018)	-0.006 (0.016)	0.063 (0.054)	0.036 (0.027)
Age	0.024*** (0.005)	0.003 (0.005)	0.036*** (0.009)	0.016*** (0.005)	0.072*** (0.018)	0.036*** (0.009)
Age squared	-0.000*** (0.000)	-0.000* (0.000)	-0.000*** (0.000)	-0.000*** (0.000)	-0.001*** (0.000)	-0.000*** (0.000)
Male	0.149*** (0.026)	0.027 (0.030)	-0.170** (0.066)	0.101*** (0.037)	0.092 (0.144)	0.089 (0.071)
Married	0.010 (0.036)	-0.052 (0.037)	-0.043 (0.060)	-0.048 (0.049)	-0.068 (0.156)	-0.035 (0.079)
Single	-0.048 (0.049)	-0.004 (0.051)	-0.054 (0.064)	-0.217*** (0.061)	-0.291** (0.142)	-0.144* (0.075)
Primary	-0.266*** (0.046)	0.024 (0.048)	-0.292*** (0.087)	-0.346*** (0.071)	-0.983*** (0.180)	-0.490*** (0.089)
Secondary	-0.190*** (0.040)	-0.017 (0.039)	-0.254*** (0.073)	-0.281*** (0.056)	-0.770*** (0.165)	-0.378*** (0.081)
Unemployed	0.005 (0.065)	0.046 (0.063)	-0.031 (0.101)	-0.032 (0.060)	-0.121 (0.236)	-0.058 (0.113)
Out of labor force	-0.032 (0.032)	-0.014 (0.025)	0.092 (0.059)	0.009 (0.035)	0.060 (0.117)	0.020 (0.056)
Big city	-0.018 (0.028)	-0.008 (0.038)	0.042 (0.051)	0.060 (0.058)	0.150 (0.120)	0.068 (0.061)
Small city	-0.049 (0.034)	-0.002 (0.019)	0.018 (0.045)	0.057 (0.044)	-0.007 (0.124)	-0.007 (0.061)
First generation	-0.106 (0.071)	0.060 (0.078)	-0.269*** (0.084)	-0.112 (0.074)	-0.596*** (0.191)	-0.270** (0.103)
Observations	3724	4165	4364	4086	3404	3404
R-squared	0.15	0.13	0.15	0.11	0.17	0.17

Notes: [1] Each regression controls for country and region fixed effects. [2] Standard errors are clustered at the country of origin level, \*\*\* significant at 1%, \*\* significant at 5%, \* significant at 10%. [3] *Trust* is the average, calculated at the country level from the World Value Survey, of the following question: “Generally speaking would you say that most people can be trusted (1) or that you can’t be too careful in dealing with people (0)? [4] The cheating variables are the answer to the following questions: “how often, if ever, has each of these things happened to you in the last five years? A bank or insurance company failed to offer you the best deal you were entitled to; you were sold something second-hand that quickly proved to be faulty; you were sold food that was packed to conceal the worse bits; A plumber, builder, car mechanic or other repair person overcharged you or did unnecessary work” The answer could take values Never (1), once (2), twice (3), 3 or 4 time (4), 5 times or more (5). [5] First generation immigrants are defined as individuals born in countries different than the country of residence; Second generation immigrants are individual born in their country of residence and whose fathers were born abroad. [5] The excluded group for education are people with college or more; the excluded group for marital status are divorced or widows; the excluded group for labor status are people employed; the excluded group for city size are people living in a country village, a farm or a home in the countryside.

**Table 8**  
**Trust and Cheating, First Generation Immigrants, by Year of Arrival**

	(1) Bank insurance	(2) Second hand things	(3) Food	(4) Plumber, builder, mechanic, repairer	(5) Times being cheated (sum)	(6) Being cheated (principal component)
Trust: new arrivals	0.663* (0.381)	0.292 (0.279)	0.473 (0.444)	0.770** (0.332)	2.022* (1.056)	1.056* (0.534)
Trust: old arrivals	0.206 (0.189)	0.114 (0.196)	0.425 (0.294)	0.123 (0.312)	1.190 (0.810)	0.578 (0.411)
Risk tolerance	0.055*** (0.014)	0.030* (0.015)	0.013 (0.018)	0.002 (0.019)	0.101* (0.058)	0.056* (0.030)
Age	0.021** (0.008)	-0.005 (0.007)	0.038*** (0.012)	0.008 (0.009)	0.057** (0.023)	0.028** (0.012)
Age squared	-0.000*** (0.000)	-0.000 (0.000)	-0.000*** (0.000)	-0.000* (0.000)	-0.001*** (0.000)	-0.000*** (0.000)
Male	0.112** (0.043)	-0.034 (0.054)	-0.170** (0.074)	0.075 (0.046)	-0.055 (0.175)	0.005 (0.090)
Married	-0.011 (0.050)	-0.011 (0.034)	-0.012 (0.086)	-0.082 (0.068)	-0.019 (0.218)	-0.010 (0.102)
Single	-0.024 (0.071)	0.066 (0.071)	-0.023 (0.111)	-0.287*** (0.102)	-0.030 (0.246)	0.003 (0.127)
Primary	-0.101 (0.106)	-0.001 (0.064)	-0.198* (0.098)	-0.279*** (0.097)	-0.607* (0.307)	-0.304* (0.158)
Secondary	-0.089 (0.084)	-0.062 (0.048)	-0.107 (0.067)	-0.203** (0.083)	-0.401 (0.254)	-0.212 (0.131)
Unemployed	0.044 (0.089)	-0.062 (0.088)	-0.114 (0.148)	-0.110 (0.109)	-0.758* (0.390)	-0.382* (0.195)
Out of labor force	-0.044 (0.058)	-0.027 (0.045)	0.084 (0.077)	-0.040 (0.050)	-0.039 (0.188)	-0.029 (0.096)
Big city	0.027 (0.055)	-0.020 (0.038)	0.008 (0.095)	-0.015 (0.094)	0.132 (0.160)	0.062 (0.082)
Small city	-0.023 (0.059)	0.006 (0.040)	-0.020 (0.045)	0.004 (0.059)	0.017 (0.142)	0.011 (0.078)
Old arrivals	0.121 (0.144)	0.121 (0.100)	0.061 (0.138)	0.218* (0.115)	0.378 (0.358)	0.201 (0.186)
Observations	1816	2035	2122	2004	1655	1655
R-squared	0.18	0.19	0.18	0.16	0.22	0.22

Notes: [1] Each regression controls for country and region fixed effects. [2] Standard errors are clustered at the country or origin level, \*\*\* significant at 1%, \*\* significant at 5%, \* significant at 10%. [3] *Trust* is the average, calculated at the country level from the World Value Survey, of the following question: “Generally speaking would you say that most people can be trusted (1) or that you can’t be too careful in dealing with people (0)? [4] The cheating variables are the answer to the following questions: “how often, if ever, has each of these things happened to you in the last five years? A bank or insurance company failed to offer you the best deal you were entitled to; you were sold something second-hand that quickly proved to be faulty; you were sold food that was packed to conceal the worse bits; A plumber, builder, car mechanic or other repair person overcharged you or did unnecessary work” The answer could take values Never (1), once (2), twice (3), 3 or 4 time (4) 5 times or more (5). [5] Old arrivals are first generation immigrants arrived in the country more than 20 years ago; new arrivals are first generation immigrants arrived in the country up to 20 years ago. [6] The excluded group for education are people with college or more; the excluded group for marital status are divorced or widows; the excluded group for labor status are people employed; the excluded group for city size are people living in a country village, a farm or a home in the countryside.

**Table 9****The effect of own trustworthiness on trust beliefs**

## A. OLS estimates of expected trustworthiness on own initial trustworthiness

	Rounds 1-3 Expected trustworthiness	Rounds 4-6 Expected trustworthiness	Rounds 7-9 Expected trustworthiness	Rounds 10-12 Expected trustworthiness
Own initial trustworthiness	0.744*** (0.0419)	0.542*** (0.0652)	0.475*** (0.0748)	0.452*** (0.0766)
Constant	0.0848*** (0.0161)	0.106*** (0.0232)	0.0763*** (0.0264)	0.0653** (0.0246)
Observations	276	208	171	171
R-squared	0.586	0.312	0.261	0.249

## B. OLS estimate of initial trustworthiness on “good values”

	Initial trustworthiness
Good Values	0.169* (0.0928)
Constant	0.211*** (0.0597)
Observations	83
R-squared	0.039

C. OLS estimates of expected trustworthiness on good values

	Rounds 1-3 Expected trustworthiness	Rounds 4-6 Expected trustworthiness	Rounds 7-9 Expected trustworthiness	Rounds 10-12 Expected trustworthiness
Good Values	0.122** (0.0588)	0.125* (0.0662)	0.122* (0.0725)	0.0515 (0.0824)
Constant	0.246*** (0.0376)	0.197*** (0.0434)	0.143*** (0.0448)	0.171*** (0.0531)
Observations	339	262	216	216
R-squared	0.025	0.027	0.027	0.004

Notes: [1] Robust standard errors, clustered by participant, are reported in parentheses, \*\* significant at 5%, \* significant at 10%. [2] Clustering by subject is appropriate because there are multiple observations for each subject due to the multiple-round experimental design. [3] Clustering by session does not change any of the significance levels in panels A and B. In panel C, clustering by session reduces the significance of the coefficient on good values in column 1 to the 10% level ( $p=0.061$ ), and increases the p-values of the “good values” in columns 2 and 3 to  $p=0.198$  and  $0.125$ , respectively. [4] The numbers of observations falls in later rounds because some sessions, due to time constraints, contained fewer than 12 rounds. [5] The number of observations falls when including our “good values” measure, because some participants did not complete the survey. [6] *Initial own trustworthiness* is the average proportion of money received that a subject would return---averaged over each possible amount that could be received---measured the first time the subject was assigned the role of receiver. To minimize contamination of this measure of trustworthiness by learning, while still maintaining a reasonable number of observations, all regressions using this measure only include subjects who were an entrepreneur for the first time in one of the first two rounds. [7] *Good Values* is the average of two measures obtained from a survey that subjects completed either a week prior or a week after their experimental session occurred: i) the emphasis, on a scale from 0 to 10, that the subject’s parents placed on being a model citizen as a value during their upbringing; and, ii) on the same scale, the emphasis their parents placed on group loyalty. We then divide the resulting average by 10 to put this measure on a scale comparable to beliefs (0 to 1). [8] *Expected Trustworthiness* is the average proportion each subject expected entrepreneurs to return within a particular round. Beliefs were elicited in an incentive-compatible manner for each possible investment level; the variable used is the average of these beliefs, for each subject, over each possible amount a receiver could receive. Beliefs were elicited regardless of the role the subject played in a particular round; if the subject was currently a receiver, they were instructed to exclude their own action from the calculation, and remunerated on this basis as well.

**Table 10**  
**Trust belief errors and economic performance in the on-line experiment**  
 OLS estimates of sender's earnings on errors in trust beliefs

	(1)	(2)	(3)
Belief Errors	1.898 (1.595)	2.196 (1.577)	2.196** (0.742)
Belief Errors Squared	-24.061*** (7.353)	-23.360*** (7.945)	-23.360** (4.798)
Constant	11.465*** (0.356)	10.995*** (0.639)	10.995*** (0.118)
Session Fixed Effects?	No	Yes	Yes
Session-Clustered Std Errors?	No	No	Yes
Observations	122	122	122
R-squared	0.05	0.07	0.07

Notes: [1] Robust standard errors are in parentheses, \*\*\* significant at 1%, \*\* significant at 5%, \* significant at 10%. [2] Belief errors are defined by the difference between a participant's estimate of the proportion of money received that a receiver will return and the actual average return proportion within each session, averaged over each possible amount a receiver could receive. This value excludes the participant's own action in the role of receiver. This yields a number that ranges from -1 to 1 for each participant.




**Table 11**  
**Earnings by trust belief categories in the on-line experiment**  
 OLS estimates of sender's earnings on dummies for trust beliefs categories


	(1)	(2)	(3)
Accurate Estimators	1.860*** (0.663)	1.773*** (0.657)	1.773** (0.500)
Over-estimator	0.311 (0.706)	0.324 (0.681)	0.324 (0.352)
Constant	9.930*** (0.525)	9.554*** (0.603)	9.554*** (0.135)
Session Fixed Effects?	No	Yes	Yes
Session-Clustered Std Errors?	No	No	Yes
Observations	122	122	122
R-squared	0.07	0.09	0.09

Notes: [1] Robust standard errors are in parentheses, \*\*\* significant at 1%, \*\* significant at 5%, \* significant at 10%. [2] Dependent variable is sender's earnings in euros. [3] The excluded category is "under-estimators." [4] Belief error categories are defined as follows: "Accurate Estimators" had an average belief error within the interval [-0.1, 0.1]; "Over-estimators" had an average belief error in the interval (0.1,1]; "Under-estimators" had an average belief error in the interval [-1,-0.1). [5] We also considered wider and narrower intervals separating the three categories, using [-0.15, 0.15] and [-.05, 0.05] to define accurate estimators. This did not change anything qualitatively; [6] Another specification used the 33<sup>rd</sup> and 66<sup>th</sup> percentiles of the error distribution in the data to separate the three categories. This did not change the results.

Figure 1  
Trust beliefs: density functions by country



Most people can be trusted or you can't be too careful




Graphs by Country


Figure 2  
The trust-income relation



Figure 3  
The empirical relationship between income and trust



**Figure 4**  
**The relationship between trust and income, by level of trust**



**Figure 5**  
**The Relationship between Trust and Income, by Level of Education and Age**

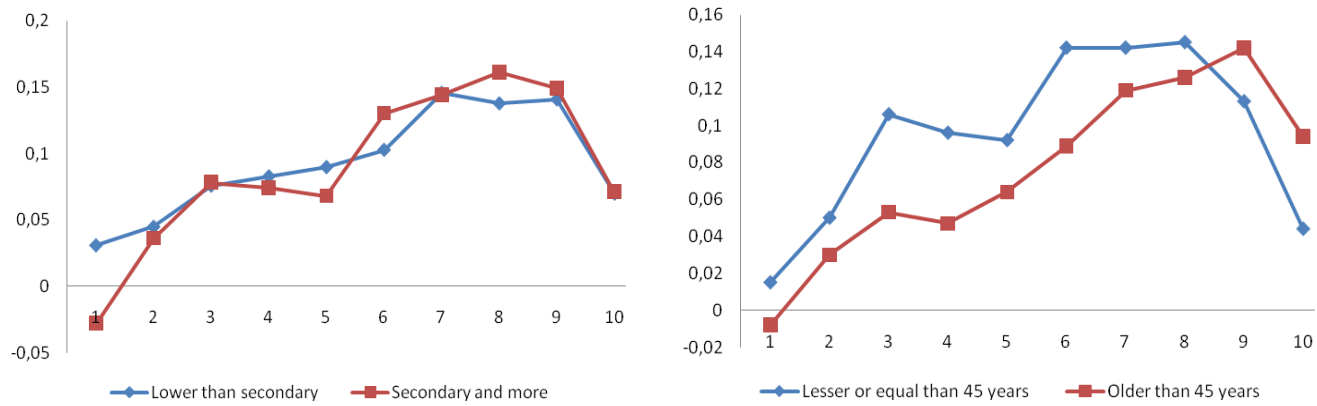




Figure 6  
Number of times being cheated




**Figure 7**  
**Heterogeneity in trust beliefs and own trustworthiness**

**A. Trust beliefs**



**B. Own initial trustworthiness**



**Figure 8**  
**Trust belief errors and performance in the on-line experiment**

