

THE RIGHT AMOUNT OF TRUST

Jeffrey V. Butler
Louisiana State University

Paola Giuliano
UCLA

Luigi Guiso
Einaudi Institute for Economics and
Finance (EIEF)

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. Our interpretation is that highly trusting individuals tend 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. However, individuals with overly pessimistic beliefs 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 forgoing college. Our findings hold in large-scale international survey data, as well as inside a country with high-quality institutions, and are also supported by experimental findings. (JEL: A1, A12, D1, O15, Z1)

1. Introduction

More than 35 years ago, Kenneth Arrow (1972), recognizing the pervasiveness of mutual trust in commercial and noncommercial 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

The editor in charge of this paper was Nicola Gennaioli.

Acknowledgments: We thank Yann Algan, Oriana Bandiera, Sascha Becker, Larry Blume, Francesco Caselli, Nicola Gennaioli, Bhajan Grewal, Erzo Luttmer, Bo Rothstein, Bernard Salanié, Guido Tabellini, Romain Wacziarg, Georg Weizsäcker, and two anonymous referees, as well as participants at various conferences and seminars. 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. Luigi Guiso and Paola Giuliano are grateful for financial support from the EUI and a UCLA-CIBER grant, respectively. Giuliano is a Research Associate at NBER and a Research Fellow at CEPR and IZA. Guiso is a Research Fellow at CEPR.

E-mail: jeffreybutler@lsu.edu (Butler); paola.giuliano@anderson.ucla.edu (Giuliano); luigi.guiso@eief.it (Guiso)




FIGURE 1. Trust beliefs: density functions by country.

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.¹ 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 quite heterogeneous across individuals. Figure 1 shows the distribution of trust for each of the countries surveyed in the European Social Survey (ESS). Here, trust is the belief about how much a generic person should be trusted, measured on a scale between 0 and 10: zero means no trust at all, while 10 means others can be fully trusted.² If this question accurately measures individuals' beliefs about the average trustworthiness of each country—a single number—respondents cannot all be simultaneously right.³ Some must have overly pessimistic beliefs, while others must have beliefs that are too optimistic. Individuals with beliefs in the tails of

1. Trust has been shown to be strongly correlated with GDP per capita and GDP growth (Knack and Keefer 1996; Knack and Zak 2001; Tabellini 2008; Algan and Cahuc 2010); with the organization of firms across countries (Bloom et al. 2009) and their ability to grow large (La Porta et al. 1997); with the size of a country's stock market (Guiso et al. 2008a); with regulation (Aghion et al. 2010); and with cross-country trade patterns (Guiso et al. 2009).

2. See Section 3 for the exact wording of the question in the ESS.

3. An alternative story is that answers reflect the trustworthiness of only those people with whom respondents interact. Because these subpopulations are likely heterogeneous, all respondents *can* be simultaneously correct. We take this concern seriously and, after developing a theoretical model, at the end of Section 2 we argue that this would work *against* our main finding.

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 overinvest in others and will be cheated more frequently, hampering their economic outcomes. Hence, at the individual level, the relationship between trust and economic performance is hump-shaped. There exists an intermediate level of trust—the “right amount of trust”—that maximizes individual income. This level of income, and trust, will be attained by individuals whose beliefs are in line with the average trustworthiness of the population they interact with.

We test the relationship between trust and income using data from the ESS. Because the survey measures the *intensity* of individuals’ trust beliefs on a scale from 0 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 eleven 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 7. Beyond a trust level of 7, income declines. The decline is initially small—however, income falls precipitously moving from the trust level of 9 to the highest trust level. This is consistent with the idea that some people make trust mistakes. The magnitude of the income cost of these mistakes is also economically relevant. On average, the income of individuals with the lowest level of trust is 14% lower than that of individuals with a level of trust corresponding to peak income. This difference is of the same order of magnitude as the income premium in our data associated with obtaining a college degree. Those who express the highest level of trust (10) make an income that is 7.3% lower than the income at the peak. Accordingly, the cost of deviating from the right amount of trust can be substantial.

While this result cannot be explained by standard forms of reverse causality (if income reverse causes trust to increase, it can explain the rising portion of the relationship but not the decreasing one; and vice versa if it causes trust to fall), there are three main issues with the pooled OLS analysis done with the ESS. The first is that the identification of the hump-shaped relationship between income and trust mainly comes from the drop in income for individuals whose level of trust is equal to 10. This raises the concern that the results could be driven by a small fraction of the population with specific characteristics. We show this is not the case in three ways. First of all, we document that individuals reporting a trust level of 10 do not differ significantly in terms of observables from the median person in their country. Secondly, we restrict attention to countries in the ESS with a relatively low average level of trust, where the identification does not come from individuals reporting 10, and we show that income is still hump-shaped in trust in these countries. Thirdly, we gather additional data on Sweden, a country in which there is a substantial fraction of observations in the upper tail of the trust distribution, and we show that the hump-shaped relationship still holds there.

The second issue with the pooled regressions with the ESS is that there might be different groups (possibly different by country), all of which have a different trust–income relationship, perhaps not necessarily inverted-U-shaped. The inverted-U-shaped relationship could then simply be the result of composition effects, obtained by pooling different groups from different countries. The Swedish data set allows us to address this concern by splitting the sample according to a large number of observable characteristics. We find that the relationship between trust and income is hump-shaped in a wide variety of subgroups.

The third concern is heterogeneity: when pooling all the countries in the ESS, we impose a single income–trust relationship and thus the same income-maximizing trust for all individuals in the sample. The results using the Swedish data are unlikely to be subject to this criticism. In Sweden there is little heterogeneity in trust by geographical regions, implying that people interact with groups of similar trustworthiness.

Beyond documenting the trust–income relationship, the ESS allows us to study one of the mechanisms through which trust beliefs affect economic performance: exposure to the risk of being cheated. The survey asks individuals whether, over the past five years, they have been “cheated” 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. We isolate the causal effect of trust on the probability of being cheated with an instrumental variable (IV) 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 remainder of the paper is organized as follows. In Section 2, we present a simple framework that predicts a hump-shaped relationship between individual trust and performance. In Section 3, we describe the survey data and present basic descriptive results from our estimation of the trust–performance relationship coming from the ESS, complemented with evidence from a large survey for Sweden. In Section 4, we estimate the effect of trust on the frequency with which one is cheated. Section 5 concludes.

2. Individual Trust and Economic Performance: A Simple Model

In this section, we present a simple framework to motivate the empirical relationship between income and trust. 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 coauthors 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 a cheater. A fraction $1 - \pi$ 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.⁴

Investors do not know the real fraction of cheaters in the population, but they do hold heterogeneous trust beliefs, τ , distributed on the unit interval $[0, 1]$. Given the presence of these (possibly incorrect) beliefs, an investor maximizes a perceived utility given by

$$\begin{aligned} \max_S \quad & Y(S) = E - S + \tau \gamma f(S) \\ \text{subj. to:} \quad & S \leq E. \end{aligned}$$

Let S_τ^* denote the optimal amount invested when beliefs about others' trustworthiness are equal to τ and let $Y(S_\tau^*) = E - S_\tau^* + \pi \gamma f(S_\tau^*)$ be the true expected income (i.e., evaluated using actual choices and true trustworthiness π). It is easy to show that whenever perceived trust is different from the truth, the true expected income, $Y(S_\tau^*)$, is a hump-shaped function of the investor's trust beliefs.⁵ This function attains its maximum when perceived trust, τ , equals the true share π (Figure 2). We will often refer to the level of trust beliefs that maximizes income as the "right amount of trust".

The model has two empirical implications. (i) The relationship between individual economic performance and trust beliefs is hump-shaped. Investors with a very low level of trust, by underinvesting, lose little if cheated but give up profitable opportunities by retaining much of their endowment (the latter effect far exceeds the former). However, investors with a very high level of trust also do worse than those close to the average trustworthiness of the population: by investing a lot in the productive venture, which

4. 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 $\pi \gamma f'(0) > 1$: 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.

5. When an individual's trust beliefs are τ , his average realized income is $Y(S_\tau^*) = E - S_\tau^* + \pi \gamma f(S_\tau^*)$ where S_τ^* is such that $\tau \gamma f'(S_\tau^*) \equiv 1$ from the first-order condition of the individual's maximization problem. This also implies

$$\frac{\partial S_\tau^*}{\partial \tau} = -\frac{f'(S_\tau^*)}{\tau f''(S_\tau^*)} > 0.$$

Differentiating $Y(S_\tau^*)$ with respect to the level of trust, τ , yields

$$\frac{\partial Y(S_\tau^*)}{\partial \tau} = \left[\frac{\pi}{\tau} - 1 \right] \frac{\partial S_\tau^*}{\partial \tau}.$$

It follows that $\partial Y(S_\tau^*) / \partial \tau = 0$ when $\tau = \pi$ and $\partial Y(S_\tau^*) / \partial \tau \geq 0$ when $\tau \geq \pi$, implying that $Y(S_\tau^*)$ is hump-shaped in τ and achieves a maximum when $\tau = \pi$.




FIGURE 2. The trust–income relation.

can potentially raise their income, they also grant partners more trust than they deserve and lose a lot when cheated (with the latter effect dominating the former.)

(ii) Furthermore, if individuals face pools of partners which differ in their trustworthiness, then observed individual performance, $Y(S_{\tau}^*)$, will, ceteris paribus, peak at a higher level of trust for individuals facing more trustworthy pools. So, for instance, one would expect income to peak at a higher level of trust in more highly trustworthy countries.⁶

6. In this simple model the channel through which trust beliefs and individual performance are related is systematic errors in beliefs. If the heterogeneity in trust reflects correctly anticipated heterogeneity in the trustworthiness of the population instead of, at least to some extent, trust mistakes, then the individual trust–income relationship should be monotonically increasing in trust. Individuals matched with less trustworthy partners correctly foresee this, express lower trust, trade less, create lower surplus, and make less income than individuals matched with more trustworthy groups. Suppose that individuals correctly foresee the trustworthiness of the group they are matched with and groups differ in trustworthiness. From $Y(S_{\tau=\pi}^*) = E - S_{\tau=\pi}^* + \pi \gamma f(S_{\tau=\pi}^*)$ differentiating with respect to τ we have $\partial Y / \partial \tau = \gamma f'(S_{\tau=\pi}^*) > 0$, since

$$-\frac{\partial S_{\tau=\pi}^*}{\partial \tau} + \pi \gamma \frac{\partial f(S_{\tau=\pi}^*)}{\partial S_{\tau=\pi}^*} \frac{\partial S_{\tau=\pi}^*}{\partial \tau} = 0$$

from the first-order condition. Thus, income is monotonically increasing in trust. Hence, correctly anticipated heterogeneity in the trustworthiness of the pool individuals interact with produces a monotonically increasing relation and is inconsistent with the hump-shaped relation we document empirically.

3. Data and Empirical Analysis

3.1. Data

3.1.1. The European Social Survey. To study the relationship between individual performance and trust beliefs, we rely on five waves of the ESS as a main data source. The ESS is a biennial cross-sectional survey administered in a large sample of mostly European nations, containing information on individuals' social values, cultural norms, and behavioral patterns. The survey has been conducted five times: in 2002/2003, 2004/2005, 2006/2007, 2008/2009, and 2010/2011. We pool data from the five waves, using information on 32 different countries.⁷ When studying the relationship between trust and cheating (Section 4), we use only the second wave because it is the only one containing the measures of cheating. For each wave, within each country, a representative sample of around 2,000 individuals is surveyed.⁸ Pooling data across countries and waves yields 228,609 observations. From the original sample of 228,609 we lose 65,973 observations because of missing data for income or the trust variable (28.9% of the whole sample, so that the restricted sample would be 162,636); an additional 60,338 observations (37.1% of the whole sample) are lost due to missing values in the explanatory variables. Our final reference sample for the baseline specification consists of 102,298 observations. The nonresponse rate in the income variable is unevenly distributed across countries. As long as the country nonresponse is related to some country characteristics, the country fixed effects should correct for the bias in nonresponse. We do, however, treat the problem of selection on observables in nonreporting income more rigorously by estimating a Heckman selection model (see the empirical analysis for details).

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. It is this feature of the ESS that allows us to investigate whether the relationship between individuals' trust beliefs and economic performance is hump-shaped. The overall distribution of the trust measure is reported in Table A.1 of the Online Appendix (for brevity we will refer to it as the OA).

7. The list of countries together with further details about the sample, the overall survey design and our variables of interest is provided in the Online Appendix (henceforth, OA).

8. Some countries are not present in all the waves. In our regressions, sample size differs by country depending on country population and ranges from 579 in Iceland (which was surveyed only once) to 14,487 for Germany (by pooling together all of the waves).

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 household's 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.⁹ 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 defined so as to be consistent across different frequencies.¹⁰ We convert all responses to their annual equivalent. The resulting brackets range from less than 1,800 per year to above 120,000 (the largest net income allowed). In our analysis we identify each bracket with its midpoint. The last bracket for the top income is coded in a country-specific way.¹¹ Table 1, panel A, shows summary statistics for (log) income in the sample. Of course, one needs to be cautious when interpreting results obtained with this measure of income, as the measure only provides income brackets and it is obviously much less accurate than the income data that could be obtained from labor force surveys or administrative fiscal sources. The fact that our results also hold within Sweden and in a laboratory setting, where the income of the subjects is given and precisely measured (see Butler et al. 2015), should reassure the reader that it is unlikely that our results are driven by measurement error.

3.1.2. The SOM Survey. To study the relationship between individual performance and trust beliefs we also use information drawn from the SOM survey, a nationwide survey carried out in Sweden from the SOM Institute, a research center studying Society, Opinion, and Media at Gothenburg University. The survey collects information on politics, society, the use of media, public service, the environment, risks, new media technology, and leisure-time activities. Starting from 1996, individuals were also asked (like in the ESS, on a scale from 0 to 10) the extent to which they believe that in general

9. While trust may affect all types of income, certain types, such as income from capital, may perhaps 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.

10. For instance, the first income category identifies income below 40 per week or below 150 per month or below 1,800 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. See the OA for more details.

11. The average income for people whose income is higher than 120,000 has been calculated using a variety of surveys. For most of the countries in our sample, we use the European Union Statistics on Income and Living Conditions (EU-SILC); data for Switzerland and Turkey (the only two countries not covered by EU-SILC) have been obtained from the Luxembourg Income Study and the Income and Living Condition Use Survey, respectively. We also run as a robustness check interval regressions and the results are essentially the same.

TABLE 1. Descriptive statistics.

Variable	Mean	Std dev.	Variable	Mean	Std dev.
A. European Social Survey					
Log income	9.716	1.072	People helpful	4.897	2.325
Age	45.309	17.966	Treated equally	4.936	1.019
Male	0.476	0.499	Helping others	3.843	0.792
Immigrant	0.084	0.278	Choose pace at work	5.734	3.611
Married	0.552	0.497	Decide daily work	6.160	3.490
Father primary education	0.373	0.484	Influence policy decision	4.041	3.601
Unemployed	0.047	0.212	Trustworthiness_a	10.204	6.407
Out of labor force	0.519	0.500	Trustworthiness_b	16.237	9.398
Years of education	12.182	4.127	Religiosity	4.585	2.953
Big city	0.331	0.471	Left-right scale	5.121	2.184
Small city	0.306	0.461	Trust	5.097	2.489
Partner primary education	0.125	0.331	Trust legal	5.218	2.629
Mother primary education	0.412	0.492	Trust parliament	4.594	2.537
Number of household members	3.084	1.230	Trust police	5.973	2.547
Professionals	0.138	0.345	Trust politicians	3.680	2.352
Technicians	0.161	0.368	Trust political parties	3.681	2.329
Clerks	0.110	0.313	Trust United Nations	5.324	2.490
Workers	0.144	0.351	Cheated: Bank	1.508	0.905
Agricultural workers	0.039	0.194	Cheated: Second hand goods	1.359	0.780
Mechan., repairers, textile work.	0.134	0.341	Cheated: Food	2.045	1.419
Assemblers, operators, drivers	0.078	0.269	Cheated: Plumber, repairer	1.617	0.995
Labourers, elementary occ.	0.103	0.304	Cheated (sum)	2.486	2.754
Risk aversion	3.034	1.412	Cheated (princ. comp.)	0.107	1.381
Altruism	5.078	0.871			
B. SOM Survey, Sweden					
Log(income)	12.546	0.652	Women	0.502	0.500
Trust	6.464	2.246	Unemployed	0.043	0.202
Age	47.723	17.510	Out of labor force	0.351	0.477
Married	0.520	0.500	Up to high school degree	0.439	0.496
Smaller conurbation	0.235	0.424	University or higher	0.313	0.464
City or bigger conurbation	0.455	0.498	Risk aversion	3.264	2.176
Stockholm, Goterborg, Malmo	0.149	0.356			

other people can be trusted. We use all the years from 1996 to 2009. The number of observations varies by wave, ranging from 1,779 in 1996 to 4,926 in 2009. Pooling data across waves yields 47,111 observations. From the original sample of 47,111 we lose 6,291 observations because of missing data for one or more explanatory variables (the restricted sample would be 40,820); an additional 1,829 observations (roughly 4% of the whole sample) are lost due to missing values in income. Our final reference sample consists of 38,991 observations.

Measuring Trust. The SOM survey elicits trust beliefs by asking the following question: “In your opinion, to what extent can one trust people in general?” Responses can take values from 0 (“Cannot trust people in general”) to 10 (“Can trust people

in general.”) One interesting feature of the Swedish data set is the high number of individuals answering 9 or 10: 3,070 and 2,703, respectively (see Table A.2 of the OA). This makes it unlikely that the hump-shaped income–trust relationship, if present, is identified only from a small number of people with peculiar unobserved characteristics.

Measuring Performance. We measure performance using the log of household income before taxes (the definition includes pensions and study allowances). The variable is defined in brackets. As with the ESS, we assign the midpoint to each income bracket and adjust for inflation.¹² Table 1, panel B, shows summary statistics for (log) income in the sample.

3.2. Empirical Analysis

Figure 1 shows the distribution of trust beliefs by country. The figure indicates 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, Iceland, and the Netherlands—the distribution has a fat right tail and the modal level of trust is quite high and at around 7. A second group, which includes several Mediterranean and 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 5.¹³ Figure 1 also shows that there is a considerable number of observations at the two tails of the distribution which also varies substantially among countries: in high-trust countries such as Denmark, the fraction of people reporting 9 or 10 is equal to 12.0% and 6.7% of the overall sample respectively, whereas the fraction of people reporting 0 or 1 is equal to 1.0% and 0.9%, respectively. At the other extreme are countries such as Turkey, where the fraction of people reporting 9 or 10 is equal to 3.3% of the sample, whereas the fraction of people reporting 0 or 1 is equal to 48.4% overall.

We start to study the relationship between individual economic performance and individual trust by estimating the following model in the pooled sample of the countries in the ESS:

$$y_{ic} = \sum_j \alpha_j \text{Trust}_{jic} + \beta X_{ic} + \delta C + \epsilon_{ic}. \quad (1)$$

Here y_{ic} is the income (in logs) 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 ten dummies $\text{Trust}_1, \text{Trust}_2, \dots, \text{Trust}_{10}$, the excluded group being individuals reporting the lowest possible trust level of 0. This specification allows

12. There are ten income brackets for the period between 1996 and 1998, eight income brackets for the period between 1999 and 2007, and nine income brackets for the 2008–2009 period. Details are provided in the OA. We also run interval regressions with the SOM survey and the results are unchanged.

13. The (whole) sample mean trust level is around 5, with a standard deviation of 2.5 (see Table 1).

wide flexibility in estimating the relationship between income and trust imposing no parametric assumptions. Finally, to control for systematic differences in average income across countries we insert a vector of country fixed effects C . Among other things, these fixed effects capture differences in individual performance due to systematic differences in the average level of trustworthiness across countries.¹⁴ The vector X_{ic} includes years of education as well as years of education of the father as proxies for acquired and inherited human capital, respectively. It also contains a linear and quadratic term in age to capture life-cycle effects in income, dummies for gender, and marital, employment, and immigration status, as well as dummies for city size with rural areas as the excluded category. This set of controls could be correlated with trust (Alesina and La Ferrara 2002) and affect income independently. To make sure that the trust question is not capturing risk aversion, altruism, or trustworthiness, we also include proxies for these three variables in our specification.¹⁵

The baseline specification indicates that the income–trust relationship is increasing for low levels of trust, before leveling off and peaking at a trust level of 7. Beyond a trust level of 7, income declines. The decline is initially small, but income falls precipitously moving from the trust level of 9 to the highest trust level. Table 2 shows the resulting estimates and Figure 3 plots the corresponding coefficients of column (1).¹⁶ Using these estimates, those with a very low level of trust (trust = 1) earn an income that is 14% lower than that of individuals with a level of trust corresponding to peak income. This difference is of the same order of magnitude as the income premium in our data associated with obtaining a college degree. Those who express the highest level of trust (10) make an income that is 7.3% lower than the income at the peak. 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 columns (2)–(5) we report evidence about robustness to some concerns with the baseline specification. Column (2) accounts for selection in reporting income using

14. In the OA we also report estimates of (1) that include regional fixed effects. These estimates (see Table A.5) give the same results as the specification with only country fixed effects. We do not include them in the main specification because of the large number of missing observations at the regional level. We always include region fixed effects when we use data from Sweden.

15. These variables are described in the OA, p. 4 and p. 5. There is a still-unsettled debate over whether questions such as those asked by the ESS or the World Values Survey (WVS) reflect expected trustworthiness only, or a mix of beliefs and individual preferences (see Miller and Mitamura 2003). Fehr (2009) points out that the answers to trust questions like those asked in the ESS likely reflect not only individuals' beliefs about others' trustworthiness, but also individuals' preferences toward risk, and in particular towards social risk. Cox (2004) has argued 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. Finally, Glaeser et al. (2000) argue that the typical trust survey question measures trustworthiness rather than trust beliefs. For a different view on this point see Sapienza et al. (2007).

16. The controls have effects consistent with our priors: income increases with own and father's education; it is higher for male and married people, lower for the unemployed, for those out of the labor force, and for immigrants. The coefficients for the full set of controls are reported in the OA, Table A.3. Risk tolerance and trustworthiness are positively related to income, whereas altruism does not have a significant effect.

TABLE 2. The relationship between trust and income.

Dependent var.: Log (income)	(1) OLS	(2) Heckman	(3) OLS	(4) OLS	(5) OLS
Trust 1	-0.003 (0.013)	-0.002 (0.013)	0.004 (0.015)	-0.003 (0.013)	-0.004 (0.013)
Trust 2	0.044*** (0.012)	0.042*** (0.012)	0.041*** (0.014)	0.040*** (0.012)	0.042*** (0.012)
Trust 3	0.070*** (0.011)	0.071*** (0.011)	0.070*** (0.013)	0.060*** (0.011)	0.068*** (0.011)
Trust 4	0.073*** (0.011)	0.074*** (0.011)	0.059*** (0.013)	0.060*** (0.011)	0.070*** (0.011)
Trust 5	0.083*** (0.010)	0.084*** (0.010)	0.071*** (0.012)	0.069*** (0.010)	0.081*** (0.010)
Trust 6	0.117*** (0.011)	0.121*** (0.011)	0.106*** (0.012)	0.097*** (0.011)	0.114*** (0.011)
Trust 7	0.140*** (0.010)	0.140*** (0.010)	0.134*** (0.012)	0.116*** (0.010)	0.136*** (0.010)
Trust 8	0.139*** (0.011)	0.141*** (0.011)	0.128*** (0.012)	0.113*** (0.011)	0.136*** (0.011)
Trust 9	0.138*** (0.014)	0.139*** (0.014)	0.140*** (0.015)	0.115*** (0.014)	0.136*** (0.014)
Trust 10	0.067*** (0.017)	0.067*** (0.017)	0.071*** (0.021)	0.056*** (0.017)	0.066*** (0.017)
Country fixed effects	yes	yes	yes	yes	yes
Individual controls	yes	yes	yes	yes	yes
Altruism, risk aversion, trustworthiness	yes	yes	yes	yes	yes
Additional controls	no	no	yes	no	no
Trust legal system (10 dum.)	no	no	no	yes	no
Controlling for moderation	no	no	no	no	yes
Observations	102,298	96,782	64,404	100,449	102,298
R-squared	0.67		0.72	0.67	0.67
Trust peak = Trust 2 (<i>p</i> -values)	0.00	0.00	0.00	0.00	0.00
Trust peak = Trust 10 (<i>p</i> -values)	0.00	0.00	0.00	0.00	0.00

Notes: 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". Individual controls in column (1) include a quadratic in age, gender, immigrant, marital, and labor market status, years of education, education of the father, and dummies for city size. The specification also includes measures of risk aversion, altruism, and trustworthiness. The variables are described on pp. 4 and 6 of the OA. Additional controls in column (3) are a full set of age dummies, a full set of education dummies, and their interactions with country dummies, mother's and partner's education, and the number of people living at home. Column (4) includes ten dummies for trust in the legal system, whereas column (5) controls for a measure of moderation, including dummies for risk aversion.

***Significant at 1%.

a Heckman selection model.¹⁷ Accounting for selection the peak of income is at a trust level of 8 and the cost of departing from the income-maximizing trust is even

17. The exclusion restriction in the selection equation is the absolute difference between the month in which the individual is interviewed and the month in which taxes are filed in a given country. We expect that this variable affects the probability of reporting income because individuals are more likely to remember




FIGURE 3. The empirical relationship between trust and income.

higher, but the qualitative result is unchanged. Column (3) tests whether the hump-shaped effect of trust on income is robust to the inclusion of a larger set of controls. In particular, we include a full set of age dummies to better capture the nonlinearity in the age–income relationship, a full set of education dummies and their interactions with each country to account for country-specific human capital effects, measures of mother’s and partner’s education (in addition to own and father’s education), as well as the number of people living at home. The inclusion of this richer set of controls (although it substantially reduces the sample size) does not alter the results. Column (4) addresses another concern: generalized trust in people could be correlated with, and therefore pick up the effect of, trust in institutions. To allow for this possibility we include ten dummies measuring the level of trust individuals have in the legal system (also available in the survey on a scale from 0 to 10).¹⁸ The “trust in the legal system” variable does not affect the hump in the generalized trust regression. Using other measures of trust such as trust in parliament, trust in the police, trust in politicians, trust in political parties, and trust in the United Nations produces similar results (Table A.6 of OA).

Another concern is that people with extreme beliefs (not only with respect to trust) could do worse in life than people with correct beliefs. If this is the case, their

income when taxes are due. Hence, the probability of reporting should decrease with the distance between the interview and tax filing. However, this distance is unlikely to systematically affect the level of reported income. The complete estimates of the Heckman model are reported in Table A.4 of the OA.

18. The question reads as follows: “Please tell me on a score of 0–10 how much do you personally trust the legal system. 0 means you do not trust an institution at all, and 10 means you have complete trust.”

income might be low not necessarily because of trust, but because they might be overly optimistic or overly pessimistic in many other domains of life. After all, a few millennia ago Aristotle 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. To rule out the possibility 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, we allow other types of traits to affect income nonmonotonically. Column (5) reports the results with the nonlinear inclusion of risk aversion, in the OA we test the robustness to other traits including altruism, political preferences, and religiosity.¹⁹ In all cases we find that the concavity of income in trust is statistically robust to these specifications.

One additional concern with the hump-shaped relationship is that it may be the result of systematic variation in the dispersion of trust beliefs with income. 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 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. In this way, 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 (Figure A.1 in the OA). What our model instead implies, since exposure to social risk increases with how much one trusts, is a positive relationship between the variance of income and trust. We investigate this prediction using our survey data by calculating the standard deviation of income for each level of trust and for each country in the ESS. The results confirm the presence of a positive association between trust and the variance of income.²⁰

3.2.1. *Further Probing of the Hump-Shaped Relationship between Trust and Income.*

The strategy followed so far provides a first good description of the data and it is useful to rule out a number of potential confounds. In this section we provide additional results to confirm the existence of a hump-shaped relationship between trust and

19. For the exact wording of the questions see the OA. Results are reported in Table A.7.

20. We run two different regressions: one where the data are collapsed at the trust–country–wave level and country and wave fixed effects are included; and another where the data are collapsed at the trust–country level and only country fixed effects are included (see Table A.8 in the OA). Of course, if there is a lot of noise in the measurement of trust and income, we might fail to detect a relationship between the variance of trust and income because of measurement error.

income. One concern with the pooled OLS regressions in Table 2 is that the hump-shaped relationship is identified entirely from individuals who respond 10 to the trust question, who may be few in number and possibly have peculiar (unobserved) characteristics that could be responsible for the results. A second concern is that apart from the shifts in the constant captured by country fixed effects, the specification restricts the trust–income relationship to be the same—and thus generates the same right level of trust—across individuals. If individuals in our sample interact with pools of people which differ in their trustworthiness and if the costs to trust mistakes are individual specific, the level of trust that maximizes income will also differ. A third concern is that the result may reflect causality running from income to trust. Here we discuss these issues. As we show, addressing one of the concerns sometimes also sheds light on another.

We address the first issue in three ways. First, we compare observables in our sample between people reporting a trust level of 10 and people whose level of trust is equal to the median of the population in each country. We find few systematic differences in observables. Compared to individuals reporting the median level of trust, individuals reporting trust equal to 10 have similar levels of education, both own and parental. These highly trusting individuals furthermore are not less likely to be migrants, are not less likely to be married, nor are they less likely to live in urban or rural areas. They are also not systematically different in terms of risk aversion or trustworthiness. We find small differences in gender (more women report 10) and unemployment status. Since unemployment is lower for people reporting 10, however, this latter difference is unlikely to drive the decline in income observed at 10. The only significant difference we find is in terms of age, for which we fully control with the inclusion of age dummies in our specification (Table A.9 of the OA).

Second, we run separate regressions for low-, average-, and high-trust countries. Once we split the sample into groups of countries with similar distribution of trust beliefs, the identification in low-trust countries no longer comes from people answering 10, as Figure 4 shows.²¹ In these countries, income peaks at a lower level of trust and starts falling already at trust 8 and 9. It is interesting to note that the peak moves from 7 to 8 and 9 depending on the average level of trust in the population (using a quadratic specification, the peak is obtained at 6, 8, and 8.25 respectively). Importantly, this is what the model in Section (2) predicts, under the mild assumption that the average trust in a country is correlated with the average trustworthiness of its population. Thus, this evidence provides independent support for the model and at the same time shows that allowing for predictable heterogeneity in the trustworthiness of the pool of people with whom individuals interact confirms the hump-shaped relationship.

Third, we report results using the Swedish data set. This country is particularly relevant for our purposes for at least three reasons. Because of the large fraction of people in the upper tail of the trust distribution, it is less likely that such individuals

21. The figure reports the coefficients of a regression for each group of countries with all the individual controls of column (1), Table 2. In Table A.10 of the OA we report the regression results for the three groups and the quadratic specification.




FIGURE 4. The empirical relationship between trust and income in low-, average-, and high-trust countries.

have peculiar characteristics. It also allows us to document that exceeding in personal trust can be harmful even when high-quality institutions are in place and when most of the population has a high level of trust.²² Finally, because Sweden is a highly homogeneous country with average levels of trust very similar across areas (Figure A.2), individuals in the sample are likely to interact with pools of people with similar levels of trustworthiness.

Results from the Swedish sample are shown in Table 3.²³ These estimates also point to a hump-shaped relationship between income and trust: as shown at the bottom of the table we can reject the hypothesis that income at trust = 10 is not lower than at the income-maximizing level of trust. Those with a trust level of 10 have an income about 7.7 percentage points lower than the peak of income, which occurs at trust level of 9. The second column uses a quadratic specification. With this specification, the income-maximizing level of trust is also equal to 9.

The data on Sweden are also useful to rule out the possibility that results are driven by composition effects. One worry with the ESS is that there might be different groups, possibly different by country, all of which have a different trust-income relationship (e.g., some positively others negatively sloped), and that the hump-shaped relationship

22. See Figure A.2 in the OA for the distribution of trust across different regions of Sweden.

23. The controls are the same individual-level variables as in Table 2. The full specification is reported in Table A.11 of the OA.

TABLE 3. The relationship between income and trust, Sweden.

Dep. var.: log(income)	(1)	(2)
Trust 1	0.038 (0.041)	
Trust 2	0.170*** (0.032)	
Trust 3	0.205*** (0.029)	
Trust 4	0.222*** (0.029)	
Trust 5	0.210*** (0.027)	
Trust 6	0.275*** (0.028)	
Trust 7	0.295*** (0.027)	
Trust 8	0.319*** (0.027)	
Trust 9	0.337*** (0.028)	
Trust 10	0.260*** (0.028)	
Trust		0.0653*** (0.006)
Trust squared		-0.0036*** (0.000)
Income maximizing trust	9.0	9.10
Trust peak = Trust 2 (<i>p</i> -value)	0.00	
Trust peak = Trust 10 (<i>p</i> -value)	0.00	
Observations	38,991	38,991
<i>R</i> -squared	0.29	0.29

Notes: Trust is the answer to the following question: "In your opinion, to what extent can one trust people in general?" The answers go from "Cannot trust people in general" (0) to "Can trust people in general" (10). Each regression controls for a quadratic in age, gender, marital, and labor market status, rural and urban areas dummies, education, risk aversion, citizenship status, and region fixed effects.

***Significant at 1%.

we document comes primarily from composition effects when different groups from different countries are pooled in the same regression. The relatively small number of observations for each country in the ESS does not allow us to split the sample by subgroups. With the Swedish data set, however, we can split the sample according to a large number of observables and check whether the shape of the trust-income relationship differs qualitatively across subgroups. Table A.12 in the OA shows separate regressions for the following groups: young and old; married and unmarried; living in rural versus urban areas, or in different regions of Sweden. We also show differences by level of education and by types of occupation (entrepreneurs versus workers, white collar, or farmers). We also split the sample by gender, by different levels of risk aversion, and by whether respondents are Swedish citizens. Finally, we split the sample

according to whether an individual grew up in Sweden or not. Reassuringly, we find a hump-shaped trust–income relationship in all of these different subgroups.

Splitting the sample according to individual characteristics (as in the Swedish data set) or based on differences in country-level averages of trust is the cleanest way of studying the presence of heterogeneous effects, given the limitations of our data. In the working paper version of this paper (Butler et al. 2009) we tackle the issue of heterogeneity more formally. We do so by relaxing the assumption that the right amount of trust is the same for all individuals and allowing for heterogeneity in the trustworthiness of the pool of people with whom individuals interact as well as in individuals' income sensitivity to trust mistakes. We do find several interesting results: first, the trust–income relationship is hump shaped for all individuals in our sample and both sources of heterogeneity drive the relationship between trust and income. Second, on average the individual right amount of trust exceeds the trust of the average person in the country. Third, though a majority of individuals has well-calibrated beliefs, a full 10% of the people in the sample make substantial trust mistakes. Finally, for around half of the sample the income cost of trust mistakes is less than 1.6% of what they would earn if they had correct beliefs (and for 3/4 of the sample is less than 6.4%); however for about 10% of the sample the cost exceeds 13% of potential income and for half of these individuals the cost is larger than 17%.

3.2.2. Reverse Causality. When looking at the correlation between individual income and trust one may argue 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 et al. (2000), and social relations enhance trust. Insofar as this reverse causality argument is true, the rising portion of the documented trust–performance relationship 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, is unlikely to be the full driver of the relationship. To rule out the possibility of reverse causality more formally, in the working paper version we attempt an IV strategy, confirming the hump-shaped relationship found with OLS. In addition, we also show that the relationship between trust and income is hump shaped in a laboratory setting, where subjects play a trust game. In the laboratory, income is the result of performance and cannot cause participants' beliefs.²⁴

4. Trust and Cheating

Two sources of suboptimal behavior contribute to the hump-shaped relationship between income and trust. On the one hand, too little trust worsens performance

24. Since in the laboratory individuals face the same pool of opponents and are randomly matched with one of them, the results cannot be due to pool heterogeneity and sorting but must be due to incorrect beliefs.

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, by 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, we can test the second channel since the ESS provides information on how often individuals have been cheated in various domains.

4.1. Measuring Cheating Experience

The second wave of 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 five ways—never, once, twice, three or four times or, finally, five times or more—which we code with the numbers 0 to 4. Figure 5 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 nonnegligible 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 five 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.




FIGURE 5. Number of times being cheated.

4.2. Empirical Specification

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

$$Z_{ic}^d = \alpha Trust_{ic} + \beta X_{ic} + \gamma C + \xi_{ic}, \quad (2)$$

where Z_{ic}^d is a summary 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 equation (1) of Section 3, but in this specification trust is a single variable (rather than a set of dummies) taking values from 0 to 10. We use a single trust measure because we are going to instrument for trust. Furthermore, in principle the risk of being cheated should increase monotonically with trust. We control for income to capture differences in the number of transactions people engage in for a given level of trust. To address the concern that trust is simply a proxy for risk attitudes, we add the survey measure of risk tolerance as a control. All regressions include a full set of occupation dummies.²⁵ Moreover, we insert into this regression a full set of country (C) dummies to account

25. In the OA, we also report a specification in which we further control for industry dummies, establishment size, and a dummy for whether the person is responsible for supervising other employees (Table A.13).

for national differences in the fraction of cheaters, and to absorb any location-specific characteristics that may encourage or discourage cheating.²⁶

Before considering the estimates of (2) we have to confront an identification issue. Since 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 (OA, Table A.14). A solution to this reverse causality problem would be to instrument current trust with the level of inherited trust, for example using second-generation immigrants' trust. This instrument, however, is likely to be invalid in this context. To see why, notice that a negative correlation between cheating and trust is not necessarily (or at least not only) the result of learning. For instance, those mistrusting could be more likely to report being cheated because they are more attentive and thus more likely to detect cheating. This reporting error in our measures of cheating biases towards finding a negative correlation. Alternatively, those mistrusting could also be more paranoid about the possibility of being cheated. They could, for example, consider as cheating any small disappointment in an exchange/relation with a counterpart, again biasing towards finding a negative correlation between cheating and trust. IV estimates can also account for these sources of bias arising from unobserved heterogeneity. However, if they are present, inherited trust is not a valid instrument in the cheating regression. In fact, though inherited trust could be a valid instrument to correct the reverse causality induced by learning, it would still be correlated with "unobserved attentiveness to cheating" or "unobserved cheating notions" breaking its validity.

To address this endogeneity issue we need a variable that systematically affects an individual's propensity to trust others, but is unlikely to respond to shocks to being cheated or to be correlated with these sources of unobserved heterogeneity in the cheating regression. To obtain this exogenous source of variation we rely on the idea that individuals, when forming their trust beliefs, are affected by their own intrinsic trustworthiness because of "false consensus"—the tendency of individuals to extrapolate the behavior of others from their own type (Ross et al. 1977). In our context own-trustworthiness can be seen as a source of information any time an individual needs to form a trust belief. In the absence of a history of information about the reliability of a pool of people, those interacting with an unknown pool form a belief by asking themselves how they would behave in similar circumstances: since they would personally behave differently, they come up with different beliefs.²⁷ Our index of trustworthiness is a measure of how much responsibility is delegated to individuals

26. These fixed effects also take care of any variation across countries in what is considered to be cheating, and that may result in different frequencies of reported cheating across countries. Our results are robust to the inclusion of regional dummies.

27. False consensus has been shown to be a persistent phenomenon: neither providing additional information about the population of interest, nor warning individuals about 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). Butler et al. (2015) show

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 (see the OA for exact wording).²⁸

We sum 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. Because 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, for example, a mechanic or a plumber. Similarly, we do not see why being more or less trustworthy should make one more or less attentive to cheating or paranoid about it, or 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 we see 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. We do acknowledge that the validity of our instrument may not be robust to the presence of unobserved heterogeneity in skills. If delegation is correlated with a worker's skill, our instrument may be invalid as skilled workers might be better able to deal with untrustworthy individuals (which would lead to a negative bias in the IV estimates); but skilled workers might also be better able to recognize that they were indeed cheated (which would lead to a positive bias in the estimates). This concern should at least partly be dealt with by the fact that we control for education and income in the cheating regressions.

Table 4 (panel A) 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, revealing a positive effect

that measures of trustworthiness obtained from behavior in a trust game predict trust beliefs in the same experiment and that, in addition, trustworthiness can be traced back to the values instilled by parents.

28. It is worth remarking that while the "delegation" measure is likely to be a valid instrument in the cheating regression it is not in the income regression estimated in Section 3, as delegation may affect income directly.

TABLE 4. Trust and cheating.

	(1) Bank insurance	(2) Second-hand things	(3) Food	(4) Plumber, builder, mechanic, repairer	(5) Times being cheated (sum)	(6) Being cheated (principal component)
PANEL A: Second stage						
Trust	0.750*** (0.196)	0.227** (0.095)	0.589*** (0.193)	0.556*** (0.163)	2.139*** (0.557)	1.098*** (0.285)
Observations	23,350	24,923	25,338	24,740	21,930	21,930
	(1) Trust	(2) Trust	(3) Trust	(4) Trust	(5) Trust	(6) Trust
PANEL B: First stage						
Trustworthiness	0.0084*** (0.0018)	0.0078*** (0.0017)	0.0079*** (0.0017)	0.0082*** (0.0017)	0.0089*** (0.0018)	0.0089*** (0.0018)
Observations	23,350	24,923	25,338	24,740	21,930	21,930
F-stat	21.95	20.10	21.10	22.39	22.98	22.98
	(1) Bank insurance	(2) Second-hand things	(3) Food	(4) Plumber, builder, mechanic, repairer	(5) Times being cheated (sum)	(6) Being cheated (principal component)
PANEL C: Reduced form						
Trustworthiness	0.006*** (0.001)	0.002*** (0.001)	0.005*** (0.001)	0.005*** (0.001)	0.019*** (0.002)	0.010*** (0.001)
Observations	23,377	24,957	25,370	24,772	21,951	21,951
R-squared	0.09	0.08	0.11	0.06	0.12	0.12

Notes: Each regression controls for a quadratic in age, gender, immigrant, marital and labor market status, years of education, education of the father, income, dummies for city size, a measure of risk aversion, eight occupational dummies and country fixed effects. 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". 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).

*** Significant at 1%; ** significant at 5%.

of trust on the number of times an individual has been cheated. Economically, the effect of trust on exposure to cheating is substantial. Increasing trust by one standard deviation increases the number of times an individual is cheated by a bank by 1.85 (1.22 times the sample mean); the frequency of being cheated when buying second-hand goods by 0.57 (62% of the sample mean); the frequency of being cheated when buying food by 1.38 (68% of the sample mean); and the frequency of being cheated by a plumber or repairer by 1.28 (79% of the sample mean). The remaining two columns show estimates using as the dependent variable the total number of times an individual was cheated in any domain (column (5)) and the first 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.

Table 4 (panel B) shows the first-stage regression, focusing on the excluded instrument. Consistent with our identification strategy the instrument has a positive effect on the level of individual trust and is highly statistically significant (with the F -stat always above 10). The reduced-form estimates of the effect of delegation (Table 4, panel C) 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. It is worth noting that the mechanisms we investigate could be of a second order compared to the loss of investment opportunities or organizational and managerial failures. It is also worth noting that while the IV estimates do document that people who trust more are cheated more, they do not show directly that being cheated translates into lost income. Overall our results should be read as suggestive evidence for a potential mechanism driving the relationship between trust and income.

5. 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. Our results hold in large-scale survey data and inside a country with high-quality institutions.

Though both excessive trust and excessive mistrust are individually costly, the data suggest that the income 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. In contrast, excessive trust may create social surplus even if this

surplus is allocated in a way that harms the overly trusting individuals. 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.

References

- Aghion, Philippe, Yann Algan, Pierre Cahuc, and Andrei Shleifer (2010). "Regulation and Distrust." *Quarterly Journal of Economics*, 125, 1015–1049.
- Algan, Yann and Pierre Cahuc (2010). "Inherited Trust and Growth." *American Economic Review*, 100(5), 2060–2092.
- Alesina, Alberto and Eliana La Ferrara (2002). "Who Trusts Others?" *Journal of Public Economics*, 85, 207–234.
- Arrow, Kenneth (1972). "Gifts and Exchanges." *Philosophy and Public Affairs*, 1, 343–362.
- Blanco, Mariana (2008). *Within-Subject Experiments on Other-Regarding Preferences*. Ph.D. thesis, Royal Holloway, University of London.
- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2009). "The Organization of Firms across Countries." CEPR Discussion Paper No. 937.
- Butler, Jeffrey V., Paola Giuliano, and Luigi Guiso (2009). "The Right Amount of Trust." NBER Working Paper No. 15344.
- Butler, Jeffrey V., Paola Giuliano, and Luigi Guiso (2015). "Trust, Values and False Consensus." *International Economic Review*, 56, 889–915.
- Cox, James C. (2004). "How to Identify Trust and Reciprocity." *Games and Economic Behavior*, 46, 260–281.
- Dixit, Avinash (2003). "Trade Expansion and Contract Enforcement." *Journal of Political Economy*, 111, 1293–1317.
- Fehr, Ernst (2009). "On the Economics and Biology of Trust." *Journal of the European Economic Association*, 7, 235–266.
- Glaeser, Edward, David Laibson, José A. Scheinkman and Christine L. Soutter (2000). "Measuring Trust." *Quarterly Journal of Economics* 115, 811–846.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2008a). "Trusting the Stock Market." *Journal of Finance*, 63, 2557–2600.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2008b). "Social Capital as Good Culture." *Journal of the European Economic Association*, 6, 295–320.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2009). "Cultural Biases in Economic Exchange?" *Quarterly Journal of Economics*, 124, 1095–1131.
- Kelejian, Harry H. (1974). "Random Parameters in a Simultaneous Equation Framework: Identification and Estimation." *Econometrica*, 42, 517–528.
- Knack, Stephen and Philip Keefer, (1996). "Does Social Capital Have an Economic Payoff? A Cross-Country Investigation." *Quarterly Journal of Economics*, 112, 1251–1288.
- Knack, Stephen and Paul Zak (2001). "Trust and Growth." *Economic Journal*, 111, 295–321.
- 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, 596–610.
- La Porta, Rafael, Florencio Lopez de Silanes, Andrea Shleifer, and Robert Vishny (1997). "Trust in Large Organizations." *American Economic Review*, 87(2), 333–338.
- Massey, Cade and Richard H. Thaler (2006). "The Loser's Curse: Overconfidence vs. Market Efficiency in the National Football League Draft." University of Chicago Working Paper.
- Miller, Alan S. and Tomoko Mitamura (2003). "Are Surveys on Trust Trustworthy?" *Social Psychology Quarterly*, 66, 62–70.

- Ross, Lee, David Greene, and Pamela House (1977). "The False Consensus Phenomenon: An Attributional Bias in Self-Perception and Social Perception Processes." *Journal of Experimental Social Psychology*, 13, 279–301.
- Sapienza, Paola, Anna Toldra, and Luigi Zingales (2007). "Understanding Trust." NBER Working Paper No. 13387.
- Tabellini, Guido (2008). "Institutions and Culture." *Journal of the European Economic Association*, 6, 255–294.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web site:

Online Appendix

Data and do files