

December 2007

Social Capital as Good Culture

Luigi Guiso

European University Institute, EIEF & CEPR

Paola Sapienza

Northwestern University, NBER & CEPR

Luigi Zingales

University of Chicago, NBER, & CEPR

To explain the extremely long-term persistence (more than 500 years) of positive historical experiences of cooperation (Putnam 1993), we model the intergenerational transmission of priors about the trustworthiness of others. We show that this transmission tends to be biased toward excessively conservative priors. As a result, societies can be trapped in a low-trust equilibrium. In this context, a temporary shock to the return to trusting can have a permanent effect on the level of trust. We validate the model by testing its predictions on the World Values Survey data and the German Socio Economic Panel. We also present some anecdotal evidence that these priors are reflected in novels that originate in different parts of the country.

JEL: E0

Keywords: social capital, culture, trust.

ACKNOWLEDGMENT: We are grateful to Alberto Alesina, Alberto Bisin, Paola Giuliano, Eliana La Ferrara, Emir Kamenic, and Guido Tabellini for helpful comments on a previous draft. Luigi Guiso thanks the European University Institute, Paola Sapienza the Zell Center, and Luigi Zingales the Center for Research in Security Prices (CRSP), the Stigler Center, and the Initiative on Global Markets at the University of Chicago for financial support. We thank Niels Patrick Joaquin for excellent research assistance and Janice Luce for editorial help.

Introduction

The concept of social capital has gained wide acceptance in social sciences and, most recently, in economics. Economists have used social capital to explain an impressive range of phenomena: economic growth (Knack and Keefer 1997), size of firms (La Porta et al. 1997), institution's design and performance (Djankov et al. 2003), financial development (Guiso et al. (GSZ henceforth) 2004, 2008), crime (Glaeser et al. 1995), the power of the family (Alesina and Giuliano 2007), innovation (Fountain 1997), and the spread of secondary education (Goldin and Katz 2001).

Given the acceptance that social capital has achieved, it is important to understand how it is accumulated and how it is dissipated. Putnam (1993), one of the fathers of the concept (at least as a characteristic of an entire community), conjectures that social capital can be the result of historical experiences. In particular, he attributes the large difference in social capital between the North and the South of Italy to the period of independence that Northern cities had as free city-states more than 500 years ago.

Although fascinating, at first sight this hypothesis seems hardly credible. Of all the economic and historical differences between the North and the South of Italy, how could we single out this one? And whatever the beneficial effects of the communal experience are, how could they have survived for 500 years through different foreign dominations and political regimes?

However, GSZ (2007) find strong support for Putnam's conjecture. They show that a free city-state experience in the Middle Ages has an effect on today's social capital even *within* the Northern regions. More importantly, they also show that the difference in social capital—between towns that had the characteristics to become independent and towns that did not—exists only in the North (where several of these towns did become independent) and not in the South (where the power of the Norman kingdom prevented them from doing so). This “difference in difference” approach eliminates the possibility that other factors correlated with the probability of becoming independent are responsible for today's differences in social capital.

By supporting Putnam's conjecture, these findings increase the puzzle. How can social capital explain the persistence of effects 500 years after the economic conditions

and the political institutions that generated them have disappeared? In order to provide an explanation for this persistence, we define social capital as “good” culture—in other words, a set of beliefs and values (GSZ 2006) that facilitate cooperation among the members of a community.

Following this approach, Tabellini (2007) builds an interesting model of the cultural transmission of cooperative values. He relies on and extends the value transmission framework developed by Bisin and Verdier (2000, 2001) and Bisin et al. (2004), in which parents optimally choose what values to pass onto their children but, in so doing, assess their children’s welfare in terms of their own values. In Tabellini’s model this creates a strategic complementarity between norms and behavior. If more people cooperate, then the payoff from cooperation increases and this expands the scope of cooperation. Conversely, an expansion in the scope of cooperation makes it easier for parents to transmit good values to their children.

In Tabellini’s model, the effect of any institutional change (such as the quality of law enforcement) is amplified and protracted over time as a result of cultural transmission. Most importantly, when individuals are allowed to choose their institutions through political voting, the equilibrium shows path dependence: if initial conditions are favorable, then individuals will transmit values of generalized cooperation and choose strong legal enforcement; if initial conditions are unfavorable, then individuals will opt for values of limited cooperation and limited enforcement. Interpreting the different historical experience of the North and South of Italy at the turn of the first millennium as being these initial conditions, the Tabellini (2007) model could possibly account for the persistence documented by GSZ (2007). The only problem we see with this story is that voting (indeed present in the free city-states) disappears with the end of the communal experience, more than 500 years ago, to be reintroduced (at a national level) a little more than 100 years ago. What mechanism perpetuated the values acquired during the free city-state experience during this 400-year hiatus?

In this paper we build a simple complementary model based on the cultural transmission of beliefs. Economic models are generally silent on how people acquire priors (i.e., probability distributions over events with which they have no experience). We posit that intergenerational cultural transmission plays a major role in the formation of

such priors. To analyze the possible distortions in this process, we build an overlapping-generations model where children absorb the prior from their parents and then, after experiencing the real world, transmit it (updated) to their own children. The reason why this overlapping-generations model is not identical to an infinitely living agent is that parents do not weigh future and current benefits exactly the same way children do, for instance because parents internalize more of the cost of their children's mistakes when they are still at home.

This intergenerationally transmitted prior affects each individual decision regarding whether to trust other members of the society and participate in an anonymous exchange. If the trust is well founded then an individual reaps substantial gains from trade. But if it is not, she will face a major loss. As a result, a pessimistic prior will induce individuals to withdraw from the market and not invest. This strategy does minimize losses, but it will prevent any update on the trustworthiness of the rest of society.

To protect children from costly mistakes, parents transmit conservative priors to them. From a social point of view, these priors are excessively conservative because parents do not fully incorporate the value of their children learning from experience. In this context we show that, if the net benefits of cooperation are not sufficiently high, then a society starting with diffuse priors will be trapped in an equilibrium of mistrust. Starting from this situation, a positive shock to the benefit of cooperation can permanently shift the equilibrium to a cooperative one even when the shock is temporary.

For analytical tractability, we derive the results under the extreme assumption that people who trust and trade immediately learn the true distribution. To test the robustness of our results to a more realistic assumption, we simulate the evolution of priors when each member of a generation takes one random draw from the true distribution. The results confirm our intuition. In particular, we show that even a brief (from a historical perspective) positive experience of cooperation (2–3 generations) can have permanent effects. This result could rationalize Putnam's (1993) conjecture that the differences in social capital between the North and the South of Italy could be due to the free city-state experience that ended more than five centuries ago.

Similarly, the model can rationalize the long-lasting effect of a history of good institutions even after these institutions have vanished. In the context of the model, better legal enforcement can be captured as a reduction in the cost of being cheated. Even a temporary reduction in this cost can permanently increase the level of cooperation as the good experience is transmitted across generations. This effect can also explain the long-lasting effect of bad colonial institutions (Acemoglu, Johnson & Robinson, 2001) or of legal origin (LLSV, 1998).

We then test the model's implications in two different samples. To test whether it is true that the speed at which the young learn is a direct function of the level of trust they inherited from their parents, we use the World Values Survey. We find that indeed the rate at which individuals learn over time is a function of the average level of trust they are endowed with when they are young. To test the link between parents' and children's beliefs, we use the German Socio Economic Panel. Not only are the children's beliefs positively correlated with those of their parents—as shown by Dohmen et al. (2007), but the intercept is negative suggesting that children's trust is lower than that of their parents.

All this evidence is consistent with culture playing an important role in transmitting beliefs, but there is no direct evidence of its long-term persistence. To provide some direct (albeit anecdotal) evidence of the persistent differences in beliefs about North and South of Italy, we compare the two major Italian novels of the nineteenth century: *The Betrothed* (in Italian *I Promessi Sposi*), written by a Northerner (Alessandro Manzoni), and *I Malavoglia*, written by a Southerner (Giovanni Verga). While both novels center on the role of Divine Providence in coping with the struggle of life, the view they transmit is completely different. In the *The Betrothed* the main characters overcome their adversity with the help of God and others; in *I Malavoglia*, the main characters not only fail miserably but also are damned for even trying to improve their human condition. The two novels differ in their degree of optimism about the human condition as well as in the level of trust placed in others. In the Northern *The Betrothed*, the help of others is the manifestation of the working of Divine Providence: you can trust others because you trust that God will help you, as the free city-states helped themselves fight against the Emperor. In the Southern *I Malavoglia*, each one (in fact, each family) is alone in facing his misfortunes. Lack of trust prevents cooperation

and any form of risk sharing. *I Malavoglia* is the best literary expression of that “amoral familism” that Banfield (1958) finds in his anthropological analysis of a little town in Lucania and identifies as the principal cause of the underdevelopment of the South of Italy. That this difference in beliefs was present more than 100 years ago (after 400 years of no election and foreign domination), is present today, and coincides with the different historical experience in the first half of the second millennium strongly suggests the intergenerational cultural transmission and persistence of these beliefs.

All these results indicate that, if we define social capital as *the set of beliefs and values that foster cooperation*, social capital can indeed explain long-term persistence.

The rest of the paper proceeds as follows. Section 1 presents a simple model of cultural transmissions of priors. Section 2 reports some simulations and shows how this model can account for the emergence and persistence of beliefs of cooperation even in the presence of temporary shocks. Section 3 tests some of the least obvious predictions of the model by using World Values Survey data and the German Socio Economic Panel. Section 4 presents a case study of how the different expectations about cooperation are embedded in two popular novels coming from different parts of Italy where social capital is very different. Conclusions follow.

1. Modeling social capital

To capture the persistence in beliefs, we build an overlapping-generations model of prior transmission. Each cohort lives three periods and comprises half of all individuals. As Figure 1 shows, each member of cohort t in its initial period starts as a child and acquires its prior from the parents. In the subsequent period, each child (now an adult) decides whether to invest an endowment x in a project that has the characteristics of the Berg et al. (1995) trust game. After investing, each individual updates her prior and transmits it to her children. In the third and last period, the individual (now mature) has another chance to play a trust game based on the information in her possession. At the end of this period, each individual dies.

1.1. Types

The economy has two types of agents, trustworthy (h) and nontrustworthy (nh). There is uncertainty concerning the fraction of the two types. In one environment (the “honest” environment) the trustworthy type predominates with a fraction $1 \geq q_1 > 1/2$; in the other environment (the “nonhonest” or “cheaters” environment), the trustworthy types are in a minority $0 \leq q_2 < 1/2$. Individuals have a prior on the share of trustworthy types. For the generation- t person, this prior probability distribution is as follows.

	“Honest” environment (prior probability $\hat{\pi}$)	“Cheaters” environment (prior probability $1 - \hat{\pi}$)
Share of trustworthy individuals	q_1	q_2
Share of cheaters	$1 - q_1$	$1 - q_2$

Here $\hat{\pi}$ denotes the prior probability of being in the “honest” environment.

1.2. Payoffs

As in a standard “trust game”, each individual receives an endowment x and can choose whether or not to invest it. If invested, the sum becomes Kx ($K \gg 1$), but the amount returned to the investor is determined by an independent player “the receiver”) who can send back any fraction of the amount received. Consistent with the experimental literature on the trust game (e.g., Fehr et al. 2003), we assume that there are two types of receivers: a trustworthy receiver, who sends back Rx , $K \geq R \gg 1$; and a nontrustworthy receiver, who sends back Lx , $0 < L \ll 1$ (where R and L are mnemonics for “return” and “loss”, respectively). Alternatively, the investor can keep her endowment and avoid the risk of being cheated. If she does not invest then she is left with x at the end of the period.

1.3. Learning

By investing, an individual learns more about the true distribution of trustworthy people in the population—information that she can use in subsequent decisions and that she can transmit to her children. By contrast, we assume that if an individual does not invest then she will not learn. This assumption may seem extreme, since people learn not only from direct experience but also from the experience of others. This latter channel, however, is

generally weaker and particularly so when people lack trust. If I do not trust others, then neither do I trust the information they report and, thus I cannot learn from them. Hence, it is not unreasonable to assume that nontrusting people find it difficult to learn from the outside environment (see also GSZ 2005, 2008).

For analytical tractability, in the first part of the analysis we will assume that an individual who invests will perfectly learn whether the fraction of trustworthy people in the population is q_1 or q_2 . By increasing the return to trade, this assumption biases the results in favor of more trade and learning. In the simulation in Section 3, however, we will assume that each investor gets a random draw from the distribution and then updates accordingly. This assumption will imply that different members of the same cohort will have different priors, which is why the problem becomes more difficult to approach analytically. Yet we will show that the results are qualitatively similar.

1.4. Returns

Consider first the second-period decision. If an individual has not invested in the first period, then she will not invest in the second period either because no new information was accrued. Thus, the only interesting case is the one where an individual has invested in the first period.

Because investing allows investors to learn about the true distribution of types, the expected return in the second period will depend only on R , L and the true share of trustworthy individuals. Let $A = q_1 R + (1 - q_1)L$ denote the expected return if the receiver is drawn from “honest” population and $B = q_2 R + (1 - q_2)L$ the expected return if he is drawn from the population of cheaters. We assume that $A > 1$ and $B < 1$, so the expected return is positive if the population is “honest” and negative if it is made of cheaters.¹

Given these assumptions, an individual who in the first period finds out she lives in a honest population will always invest in the second period (because she expects to obtain $A > x$). But if she finds out that she lives among cheaters then she will never invest in the second period, since $x > B$.

¹ These two assumptions imply that $q_2 > (1 - L)/(R - L) > q_1$.

Ex ante, an individual expects to observe q_1 with probability $\hat{\pi}$ and q_2 with the complementary probability. Hence, at the beginning of the first period, the expected second-period return will be $\hat{\pi}Ax + (1 - \hat{\pi})x$.

The payoff for the first period is:

$$\hat{\pi}[q_1R + (1 - q_1)L]x + (1 - \hat{\pi})[q_2R + (1 - q_2)L]x = \hat{\pi}Ax + (1 - \hat{\pi})Bx .$$

Thus, the net expected payoff from investing over the two periods will be

$$P(\hat{\pi}, R, L, q_1, q_2) = \hat{\pi}Ax + (1 - \hat{\pi})Bx + \hat{\pi}Ax + (1 - \hat{\pi})x - 2x . \quad (1)$$

where $2x$ is the lifetime value of the endowment of an individual who does not invest in either period. Clearly, in the first period an individual will invest if and only if $P(\hat{\pi}, R, L, q_1, q_2) \geq 0$.

As we show in the Appendix, there exists a threshold $m = m(R, L, q_1, q_2)$, with $0 < m < 1$ defined by $P(m, R, L, q_1, q_2) = 0$, such that it is optimal to invest if and only if the received prior is above this threshold (or $\hat{\pi} \geq m(R, L, q_1, q_2)$). Furthermore, m is decreasing with R (the amount returned by the trustworthy receiver) and q_1 and q_2 (the shares of trustworthy individuals), it is decreasing with $1 - L$ (the loss incurred when the money is sent to a nontrustworthy individual).

1.5 The parent's problem

Having modeled the investment behavior, we can now determine the optimal prior that parents want to instill in their children. This optimum obviously depends upon the parents' objective function.

In order to prevent our overlapping-generations model from unraveling into a model with infinitely living agents, we assume that parents do not weigh future and current benefits exactly the same way as children do.² One way to interpret this limitation

² The overlapping-generations structure that we impose here would not have the same properties as an infinitely lived agent even if there were no parent-child gap in preferences. The reason is that the parents, by assumption, do not take into account the superior information their children will have if the parents invest.

is that parents internalize more of the cost of their children's mistakes when they are still at home (because the parents must pay for them). Alternatively, we can interpret it as parents suffering more if their children get into trouble while the parents are still alive.³ To capture this idea, we assume that parents care only about their children's first-period utility. But when parents assess the value of first-period investment, they use their own knowledge of the probability distribution of trustworthy individuals in the population.

Let π denote the parents' subjective beliefs (derived as the posterior from their endowed prior and their learning) that the true share of good types in the population is q_1 . The first-period net expected utility of the child from investing x as perceived by his parent is:

$$P_p(\pi, R, L, q_1, q_2) = \pi Ax + (1 - \pi)Bx - x = (\pi A + (1 - \pi)B - 1)x. \quad (2)$$

The parent will be indifferent between the child investing x and not investing if her prior is such that $P_p(\pi, R, L, q_1, q_2) = 0$. The threshold for indifference is then

$$m^p = \frac{1 - B}{A - B}$$

Because parents do not discount the value of information obtained by investing in the first period, it must be that $m^p > m$. The parent will then choose the prior to transmit so as to maximize (2), and her teaching strategy will then be:

$$\begin{aligned} \hat{\pi} &\geq m^p \quad \text{if } \pi \geq m^p, \\ \hat{\pi} &< m^p \quad \text{if } \pi < m^p. \end{aligned}$$

The solution is shown in Figure 2. If a parent has a sufficiently optimistic prior ($\pi \geq m^p$), then she will transmit to her children a prior that is at least as optimistic as hers. As a

³ Some may argue that a "nasty" parent would have her children take risks in order to learn from their mistakes; the model's implications would then be opposite. But this does not seem to capture normal parental behavior.

result, her children will invest and (if the population is an honest one) will continue to transmit their optimistic priors to their children and so on.

If, instead, the parent has a sufficiently pessimistic prior ($\pi \leq m$), then she will transmit a prior that is at least as pessimistic and sometimes even more pessimistic (in the region $m < \pi < m^p$) to her children, who will choose not to invest. Since the no-investment strategy does not allow for any learning, the same pessimistic prior will be transmitted unchanged from generation to generation. This pessimism will trap society in a no-trust–no-trade equilibrium, even when the majority of individuals is trustworthy. In Bisin and Verdier (2000, 2001) and Tabellini (2007), cultural transmission of norms is motivated by parents’ preferences for having kids with similar traits, but in our model the parents instill beliefs so their kids won’t make mistakes. But these beliefs may well differ from their true ones, since parents may benefit from teaching downward-biased beliefs.

As an illustration, consider the extreme case where the share of trustworthy people is 1 in the honest population and 0 in the cheating one (i.e., $q_1 = 1, q_2 = 0$). In this case the respective thresholds are

$$m^p = \frac{1-L}{R-L} > m = \frac{1-L}{2R-(1+L)}$$

If the diffuse prior is such that it exceeds $m_p = \frac{1-L}{R-L}$ which it will when $R > 2-L$. then individuals will invest and trust from that point on. If, on the contrary, the diffuse prior is below the threshold, then no one would invest and there would be no possibility of disconfirming their prior.

1.6. Comparative static

Suppose the economy starts with a diffuse prior, attaching probability 0.5 to the two distributions, and suppose the true underlying distribution is the one with a majority of trustworthy individuals. If the economy parameters are such that this prior is below m^p , then parents will not invest and may transmit an even more conservative prior to their children, who will not trade themselves and will transmit mistrust to the subsequent generation and so on.

Marginal changes in the return to investment or to the share of trustworthy individuals will have no impact on individual decisions. For instance, if $R = 1.5, L = 0.1, q_1 = 0.8, q_2 = 0.1$, then $m = 0.633 < m^p = 0.7756$ and thus, with a diffuse prior, parents will not trade and will induce a no-trade prior in their child. An increase in R to 1.65 would not alter this equilibrium; it would only change the child and parent thresholds to 0.522 and 0.687, respectively.

However, a big shock to the benefits from trusting can change the equilibrium. If, for instance, the return obtained from the receiver R jumps to $R' > R$, so that $m^p < 1/2$ after the shock, then parents will start to teach their children an optimistic prior and the economy will emerge from the low-trust trap. Using the figures of the example, the posterior probability that the receiver is drawn from the distribution with a majority of trustworthy people becomes 1. The information so acquired will be transmitted to the next generation, which will be endowed with a prior between m^p and 1.

Suppose now that the higher level R' declines to its initial value after a few generations. If the prior received by the informed generation and transmitted to the subsequent one is sufficiently optimistic, then individuals may continue to invest even though the return to investment has reverted back to its initial value. For this to be the case, we must have $\pi(q_1 | R') > m^p(R)$; that is, the inherited prior accumulated after the big positive shock must exceed the parents' threshold when the reversal occurs.

1.7 Limitations of the model

The model makes some stark assumptions, but the results' intuition survives when we relax them. In the next section we show that our results still hold under more realistic hypotheses on the updating process.

On a different count, the model assumes only an intergenerational transmission of information. In fact, we know that children learn not only from their parents but also from their peers. As Bisin and Verdier (2000, 2001) show, however, peer learning increases the complementarity of transmitted values (or, in our case, beliefs). Hence, mistrust will persist even more when there is intragenerational contagion.

Another key assumption regards the frequency of investment possibilities. We allow people to invest only their whole endowment, and we give them only two chances in their lifetime. If we allowed infinitesimal trades (i.e., trades involving very small amounts) or infinitely repeated ones, then learning will take place immediately.

Yet there are good reasons for making our two assumptions. First, infinitesimally small trades might prove to be uninformative about the true nature of the population: punishment (be it legal or moral) is related more to the act of cheating as such than to the amount at stake. Hence, a rational cheater would cheat only when dealing with sizable trades, leaving few data points for learning. Second, important choices in life are rare and take a long time to pan out. What type of career should I undertake? Should I trust my boss or the company I work for? What type of pension fund should I invest in? What person should I marry? In these choices, we learn our mistakes only in the long term and seldom have other opportunities to try again.

Finally, in our model we assume that trustworthiness is exogenously given and is not affected by the prevailing level of trust. In reality, there could be two channels through which beliefs can affect trustworthiness. First, a receiver who knows that the sender expects him to cheat is more likely to cheat, as shown by Ruben et al. (2007). Thus, mistrust breeds mistrust. Second, social pressure will make it easier to teach children to be trustworthy (a value) when the expectation (a belief) is that most people will be trustworthy.

Both these effects will strengthen the results of the model and the persistence of the equilibrium. These effects also show the complementarity between our model and Tabellini's (2007) model. Tabellini addresses the transmission of values, and we address the transmission of beliefs. Social capital is formed by both.

1.8 Empirical implications

This simple model of beliefs transmission can account for several existing facts. It also generates some new empirical implications. First, the model can explain why the trust of second-generation Americans is correlated with the level of trust in their ancestors' country of origin and why it takes several generations for this correlation to disappear (see Rice and Feldman 1997; GSZ 2004, 2006; Tabellini 2005, 2008). Immigrants carry

their trust of origin, which they transmit to their children. Experimenting in the new environment slowly modifies the imported priors, but this process takes several generations. This implication differentiates ours from a model of cultural transmission à la Tabellini (2007), where immigrant parents immediately adapt to the new environment. Second, the model can explain why the level of trust increases with the age of the individual, as found in GSZ (2003, table 2): children inherit excessively conservative priors that they slowly liberalize with experience.

The model also has more subtle predictions about the speed of this learning, predictions that we will test in Section 4. For given values of R and L , incentives to invest and increase trust depend on the initial prior and on the share of trustworthy individuals in the population. If the prior is very low—that is, if individuals mistrust to begin with—then none would invest and thus the age–trust profile should remain relatively flat. For intermediate levels of trust we should see a steeper profile. Furthermore, since chances of building trust are stronger when the proportion of trustworthy individuals is higher, we should see a steeper learning profile in countries that have a relatively high share of trustworthy people.

The model has some additional implications for the effect of legal origin and quality of enforcement in a country. If the share of trustworthiness is higher in countries with better legal systems, then we should find that, *ceteris paribus*, the trust of second-generation Americans is higher if their ancestors come from lands with better legal systems which would have given them stronger incentives to invest and learn about the trustworthiness of others.

2. Simulations and implications

To check the robustness of the model to a more realistic updating process—one where individuals learn only by observing the outcome of their trade—we simulate this alternative in MATLAB. All simulations involve 40 generations (or about 1,100 years assuming a generational gap of 25 years); each generation is composed of 100 people. We assume that the fraction of trustworthy people can be either 0.8 or 0.2. We assume that the first-generation prior is diffuse (both events are equally likely) but show only the

simulation where the true fraction of trustworthy people is 0.8, since the other case is uninteresting.⁴

Figure 3a shows the average beliefs of each generation, when $R = 1.8$ (an annual return of 2.4% over a 25-year horizon) and $L = 0.01$. In this case, the value to which the mean of the priors converges is not unique. However, since each simulation always ends up with a mix of 1s and $m - \varepsilon$, the mean of priors does converge in each run.

In the second simulation (not reported), we change the value of R to 1.5 (an annual return of 1.64%). This R -value is sufficiently low that, with a flat prior, everyone abstains from investment from the very beginning. The mean of priors jumps from the initial 0.5 to $m - \varepsilon$ (about 0.6955 here) and remains there unless a shock occurs.

In the third simulation (Figure 3b), we change the R -value to 2. This value is sufficiently high that all family lines have priors that converge to 1. (The first simulation, where the means converged to about 0.74). These outcomes confirm the analytical results: If the return to trust is sufficiently low, an economy can be stuck in a low-trust–no-trade equilibrium forever.

We proceed to analyze the effect of a temporary shock with the goal of seeing whether a model like ours can explain the long-term persistence of a (relatively) brief historical experience of cooperation. Putnam (1993) conjectures that the differences in social capital between the North and the South of Italy could be due to the free city-state experience, which lasted about 250 years and ended more than five centuries ago. As discussed in GSZ (2007), free city-states emerged as a response to the growing need for cooperation engendered by the expansion of trade routes and the simultaneous collapse of Imperial authority.

We attempt to calibrate the free city-state shock by increasing the return to cooperation R from 1.5 to 2 (an annual return of 2.8%). The shock is introduced at generation 3 and lasts through generation 12 (i.e., 10 generations or approximately 250 years). After this period, R returns to 1.5.

⁴ We need to make an assumption about how the prior transmission rule is implemented. For concreteness we assume that, when transmitting priors to a child, parents set them to the highest level consistent with the rule. In particular, we assume that if their posterior exceeds the threshold for the teaching decision then parents transmit that exact posterior to the next generation. Otherwise, they set the next generation's prior to the maximum value at which the child would not invest: at $m - \varepsilon$, where ε is a small number.

As Figure 3c shows, this “temporary” shock is sufficient to induce almost all family lines to have an optimistic prior and always invest. Most interesting, this effect persists forever even after the shock disappears.

The permanence of these effects obviously depends on the duration and magnitude of the shock. To demonstrate this, in the fifth simulation (not shown) we perturb the economy with a smaller shock (from $R = 1.5$ to $R = 1.6$). In this case we find 36% of the family lines permanently abstaining from investment (a large difference from the almost 0% eventually abstaining in the previous example), even when the positive shock lasts 20 generations.

How long should a shock last to have permanent effects? If a shock that moves R from 1.5 to 2 lasts three generations (roughly 75 years), 19% of family lines eventually abstain from investment. If it lasts for five generations, 8% of family lines eventually drop out. If it lasts eight generations, only 2% of family lines eventually drop out. Thus, even relatively temporary shocks can have long-term effects. The fundamental reason why is that the institution transmitting the accumulated culture—the family—is itself persistent and more so than any other, as Williamson (2000) implicitly assumes. This cultural transmission is very strong and can be blocked only by removing children from their parents’ influence at an early age.⁵

3. Empirical evidence

3.1. Evidence from the World Values Survey

To test the foregoing implications we first used data from the World Values Survey (WVS), pooling three waves of data (1981–4, 1990–3, and 1995–7) and regress generalized trust—our empirical measure of the model’s individual subjective probability of being cheated—on a number of determinants including age. Generalized trust is measured by answers 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?*” The

⁵ Breaking family cultural transmission was Bismarck’s main motivation for creating compulsory schooling. A more dramatic episode is that of the “lost generation” in Australia: between 1800 and 1969, aborigine children were systematically taken from their parents, and transformed into slaves of white Australians in order to break the transmission and preservation of aborigine culture.

variable is set equal to 1 when participants report that most people can be trusted and 0 otherwise.

Table 1, first column, shows the results of the estimates of a linear probability model, where country-level dummies are inserted to account for systematic differences in the average level of trust across communities. Individual trust increases significantly with age. Moving from age 30 to age 60—roughly the gap between two generations—raises the average trust by about 5 percentage points, or 20% of the sample mean.

This result is consistent with GSZ (2003), who find that, in the countries sampled by the World Values Survey, the level of generalized trust of individuals is increasing with age. Alesina and La Ferrara (2002) and Uslaner (2007) obtain a similar result using a sample of Americans in the General Social Survey.

Column 2 reports the same estimates restricted to the sample of low-trust countries (with an average trust below the 25th percentile of the distribution) and column 3 for the sample of high-trust countries (with an average trust above the 75th percentile of the distribution). As predicted by our model, the age–trust profile is flatter in low-trust countries because, when trust is too low, there are no incentives to invest and learn and thus no updating with the received prior.

The age–trust relation is almost twice as steep in countries with a high level of trust than in countries with low average trust, and the difference is statistically significant (see the *t*-test in the last row).

We further explore the model’s implication in Table 2. Incentives to invest and revise the prior depend both on the level of trust of the young and on the difference between the true and perceived trustworthiness of others. If the gap between the current trust of the young and the trustworthiness of the counterparts is large, then there is more scope to update priors and the trust of the older generation should differ more from the trust of the younger generation. In addition, the speed at which the updating takes place should be a function of the prior that the young inherit.

To test these predictions we need a proxy for trustworthiness, which we constructed by averaging out the WVS answers about individuals’ attitudes toward some noncooperative behavior. Specifically individuals are asked “Tell me for each of the following statements whether you think it can be justified” (1=always, 10=never)”:

“Claiming government benefits to which you are not entitled”; *“Avoiding a fare on public transport”*; *“Cheating on taxes if you have a chance”*; *“Buying something you know was stolen”*; *“Accepting a bribe in the course of your duties”*. We averaged out the answers to these questions and divided the result by 10 in order to have the variable on a fraction scale. Finally, to obtain an estimate of the trust gap, we subtracted from this measure of trustworthiness the average trust of the young in the country.

We construct a proxy for the degree of updating of the prior by taking the difference between the average trust of the “young” (those not older than 30 years) and that of the “old” (older than 30) by country. We can then test the model implications by estimating the following regression:

$$Trust_{OLD} - Trust_{YOUNG} = a + b(Trust\ Gap) + c(Trust\ Gap) \times Trust_{YOUNG} .$$

Results of the estimates of this model are shown in Table 2, first column. Consistent with the model prediction, the perceived trust gap has a positive and significant effect on the degree of people’s updating in trust over time. Furthermore, the effect of the perceived trust gap is stronger in countries where the young receive a conservative prior, because they experiment and learn less.

Column 2 reaches the same conclusion by using the coefficient on age of a regression of individual trust on age and the same controls shown in Table 1. Although these results are indeed consistent with the notion that parents transmit conservative priors, one must remain cautious. First of all, the regressions of trust on trustworthiness may simply reflect that, as argued by Glaeser et al. (2000), people may tend to confuse the meaning of the two words and thus, to some extent, trustworthiness may be just an alternative measure of trust and not truly indicate incentive to trust in societies where people are more trustworthy. This worry, however, is reduced by the finding of Sapienza et al. (2007) that the World Values Survey question is a good measure of people’s expectation in the Berg et al. (1995) trust game. Some lack of correlation with the amount sent may occur because, in the context of the trust game, this decision is affected by other factors (such as altruism). Second, since we are using cross-sectional data to compare the trust of the young to that of the old, differences in trust levels across individuals of

different age may reflect cohort effects rather than a trust–age relation. A proper test of the model prediction would require longitudinal data on trust, a type of information that to our knowledge is not available.

3.2. Evidence from the German Socio Economic Panel

To address some of these concerns and to shed further light on the transmission of priors across generations, we next utilize the German Socio Economic Panel (GSEP). In the 2003 wave the GSEP collected information on the trust of parents as well as on that of children, making it possible to relate the two directly (see e.g. Dohmen et al. 2007, who find a robust positive correlation between the two). In addition, the correlation is stronger between the trust of the mother and that of the child, consistent with the idea that mothers play the larger role in educating children and thus in transmitting values and beliefs.

This evidence bears directly on the role of the family as a vehicle of cultural persistence⁶ and thus supports the essence of the model in Section 2. But our model implies also that, when transmitting priors, parents are conservative and so the prior they transmit is, on average, lower than the one they hold. Furthermore, parents with stronger beliefs about the trustworthiness of others should transmit less conservative priors. By using the GSEP we can test these implications directly. For this we rely on answers to the following question: “*What is your opinion on the following three statements? 1) On the whole one can trust people; 2) Nowadays one can't rely on anyone; 3) If one is dealing with strangers, it is better to be careful before one can trust them*” For each statement, the respondent can answer in one of four ways: “Totally agree”, “agree slightly”, “slightly disagree”, or “totally disagree”.

We have recoded the answers so as to obtain three indicators of trust. The first, which we call “Generalized trust”, is a dummy equal to 1 if the respondent answers “Totally agree” or “agree slightly” to the first question and equal to 0 otherwise; the second indicator, called “Reliance on others”, is 1 if the respondent answers “slightly

⁶ Cipriani et al. (2007) find no evidence of intergenerational transmission of values in an experiment involving a standard public good game in a group of Hispanic and African American families. However, their result may reflect the fact that the experiment’s child subjects were all still attending elementary school—too young to absorb their parents’ beliefs and norms or to permit sufficiently precise measurement of a possibly incomplete transmission.

disagree” or “totally disagree” to the second question (0 otherwise); and the third indicator, called “Careful with others”, is 1 if the respondent answers “slightly disagree” or “totally disagree” to the third question (0 otherwise). Thus, each of these measures is an indicator of trust. These questions were asked to all the sons and daughters of age 18 or higher at the time of interview and to their parents. The GSEP interviews all sons and daughters (even if they have left their parents’ home) and ensures that all respondents answer separately from one another so as to avoid introducing artificial correlations in the answers (for details about the GSEP, see Schupp and Wagner 2002). We denote variables referred to the mother with $_m$ and those referred to the father with $_f$. To obtain a measure of perceived trustworthiness of the parents we use the answers to the question: “Do you believe that most people: a) would exploit you if they had the opportunity; b) would attempt to be fair towards you?” Those parents who choose the second answer have higher expectations that other people are trustworthy. Hence, we construct an indicator called “Trustworthiness” that is equal to 1 for those parents who choose answer (b). Even this measure cannot dismiss the criticism that people confuse trust and trustworthiness, but any affect that trustworthiness has on the prior that parents teach to their children (via its effect on the trust of the parents) is controlled for because we observe the trust of the parents.

To test our model’s implications, we run the following probit regression:

$$Trust = a + bTrust_m + cTrust_f + dTrustworthiness_m + eTrustworthiness_f ,$$

where *trust* on the left-hand side is one of the three measures of the trust belief and where *trust_m* and *trust_f* are the corresponding beliefs of the parents. Our model implies that: the constant in this regression should be negative, since parents transmit conservative priors; that *b* and *c* should be positive, since children inherit the parents’ beliefs; and that *d* and *e* should be positive, since parents with stronger beliefs about the trustworthiness of others teach a less conservative prior.

Table 3 shows the results for the three measures of trust (“Generalized trust” in columns (1) and (2), “Reliance on others” in columns (3) and (4), and “Careful with others” in columns (5) and (6). In all the estimates, the parameters *b* and *c* are positive

and the response to the mother's prior is larger than to the father's (as in Dohmen et al. 2007).⁷ Most interestingly, the constant is always negative, which is consistent with parents teaching a conservative prior as predicted by our model. Finally, when added to the regression, parents' beliefs about the trustworthiness of others have a positive effect on the prior that they transmit to their children. These results do not change if we add such additional controls as stating where the family is located, the level of education of the parents and the children, and the children's gender and age (see the regressions in columns (3), (6), and (9) respectively). In particular, since we are controlling for the ages of the parents and the children, the trust gap cannot be due to a cohort effect.⁸

4. Poems, novels and the transmission of social capital

The previous evidence is consistent with beliefs and norms being transmitted from one generation to another. Even the most direct evidence relating the beliefs of children to those of their parents, in Table 3, shows that there is a cultural persistence but does not show how much culture persists. For culture to be a credible vehicle of persistence of historical episodes (such as those studied in Tabellini (2005) and GSZ (2007)), we need to show that beliefs and norms may extend well beyond the span of two generations.

The correlation between the trust of second-generation Americans with the level prevailing in their ancestors' country of origin is consistent with beliefs persisting for two generations. But as shown by Soroka et al (2003), this dependence tends to vanish after the third generation as individuals slowly adapt their beliefs, which are ultimately determined experiences in the country where they live. Thus, this evidence can document persistence of cultural beliefs for about 60–90 years. The question remains of whether social capital can persist over many centuries.

⁷ Our sample has somewhat more observations than the one used by Dohmen et al. (2007), who include only those observations in the 2003 wave that are also present in the 2004 wave of the GSEP.

⁸ One may wonder how the idea that parents instill conservative priors squares with the evidence that the younger have higher risk tolerance than the old (see e.g. Guiso and Paiella 2007). There are at least two answers to this concern. First, trust and risk tolerance are conceptually different. Trust is a prior and as such is largely driven by available information, including the prior that is transmitted by parents; risk tolerance is a preference parameter and as such may also reflect innate features that even parents may be unable to reshape. Second, and consistent with these two measures reflecting different concepts, empirically the belief component of trust (as measured, for instance, by the World Values general trust question) tends to be poorly correlated with empirical measures of risk attitudes (Sapienza et al. 2007; GSZ 2008).

Because we did not observe the beliefs of populations four or five hundred years ago, we cannot document cultural persistence directly. Instead, we try to do so by comparing the poems and novels of different cultures. If culture is the vehicle through which important historical episodes are transmitted over centuries, then this should be reflected in the literature that these cultures have produced.

The obvious country to look for this difference is Italy. As we have argued (GSZ 2004, 2007), there is a massive difference in the stock of social capital between the two parts of Italy and also in the beliefs and norms of the populations in these two areas. For instance, according to the late 1990s wave of the World Values Survey, the WVS, the fraction of people who trust others is 25% in the South but 42% in the North. Similarly, the share of people who deem it important to teach obedience to their children—a measure of the hierarchical view of society and thus of the lack of personal freedom (Tabellini 2005)—is 37% in the South and 26% in the North; the share who think it is important to teach children tolerance and respect for others is 50% in the South and 60% in the North. If this lack of faith vis-à-vis other people is deeply rooted in history, then we should be able to detect it in past literature.

Toward this end we take the two most important Italian novels of the 19th century, The Betrothed by Alessandro Manzoni (a Northerner) and I Malavoglia by Giovanni Verga (a Southerner).

The Betrothed (published 1827) is cast in the 17th century, in Lombardy, between Como and Milan during the terrible, oppressive years under Spanish rule. The setting is a veiled attack on Austria, which controlled the region at the time of writing, and reflects the sentiment of independence that had carried over since the communal movement—as frequent references to the Republic of Venice remind the book's readers.

The novel tells the story of a young couple of humble origin (Renzo and Lucia) who fight against a cruel, powerful, and despicable nobleman (Don Rodrigo) who wants to prevent the young couple from marrying because he is interested in Lucia. While the difference in power between the couple and Don Rodrigo is enormous, Renzo and Lucia have the confidence to fight Don Rodrigo and, in the end, prevail.

The novel is marked by a strong sense of optimism: optimism about the future and optimism inspired by the confidence in the help of others. As Renzo and Lucia conclude,

"troubles often come, yes, because we've given us a cause; but that the most cautious and innocent conduct isn't enough to keep them away; and ... when they come, with guilt or without guilt, the trust in God sweetens them, and makes them useful for a better life."

Of a completely different nature are the attitudes and values conveyed in *I Malavoglia*, first printed in 1881. This novel describes the story of a family of fishermen (the Malavoglia) who live in Aci Trezza, a small Sicilian village near Catania. One day this family decides to use their boat to engage in a trade of fava beans. Not able to rely on any form of cooperation, the Malavoglia can neither share the enterprise's risk nor insure it. As a result, they are completely vulnerable to nature. But nature is adverse. They lose everything in a storm: the boat, the freight, and ultimately their house that had been offered as collateral to a loan shark before they could buy the fava beans. Lack of trust in people and the State, poverty, and resignation all feed the pessimism that permeates the story.

The Malavoglia's attitude toward other people is the same attitude that years later Banfield (1958) would label "amoral familism". This phrase was used to describe a culture of mistrust toward the community at large, perceived as inimical, and reliance on closed but safe family links.

Although we are unable to trace these cultural values back to the Middle Ages, it is interesting that they were well rooted more than 100 years ago in a manner consistent with the historical experience Putnam describes.

5. Conclusions

Social capital has recently gained wider acceptance in economics. Social capital's success is due in part to its remarkable correlation with economic performance across countries (Knack and Keefer 1996) and regions (Tabellini 2005). Recent empirical work (Tabellini 2005, 2008; GSZ 2007) also suggests that these correlations may reflect a causal link and explain the persistence of performance and good institutions.

In this paper we have built a model of intergenerational transmission of beliefs that can explain long-term persistence. We then argued that, together with Tabellini

(2007), our model can be used to introduce a new definition of social capital as *the set of beliefs and values that foster cooperation*. This definition makes social capital easy to measure and to incorporate into standard economic models, which we hope will overcome the legitimate skepticism of many economists for this concept.

What this paper does not address—but future research should—is the policy dimension. Once a community (be it a town or a country) is trapped in a low-trust equilibrium, what can be done to rescue it? This is next in our agenda.

Appendix

Proposition 1

Assume $A = q_1R + (1 - q_1)L > 1 > B = q_2R + (1 - q_2)L$. This ensures that if one is certain that the true population has a majority of trustworthy people then it is profitable to invest. There exists a threshold $m = m(R, L, q_1, q_2)$, with $0 < m < 1$, defined by $P(m, R, L, q_1, q_2) = 0$ such that it is optimal to invest if $\hat{\pi} \geq m(R, L, q_1, q_2)$ but it is optimal to abstain otherwise.

Proof: We are looking for a value of $\hat{\pi} = m(R, L, q_1, q_2)$ such that

$$P(\hat{\pi}, R, L, q_1, q_2) = 2\hat{\pi}Ax + (1 - \hat{\pi})Bx + (1 - \hat{\pi})x - 2x = 0$$

The threshold m is then equal to

$$m = \frac{1 - B}{2A - (B + 1)},$$

and $0 < m < 1$ because $A > 1 > B$.

Proposition 2

The value of m is decreasing with the amount returned by the trustworthy receiver, R , and with the shares of trustworthy individuals, q_1 and q_2 ; it is decreasing with $1 - L$, the loss incurred when the money is sent to a non-trustworthy individual.

Proof: The proof follows immediately by taking first derivatives of m with respect to R , L , q_1 , and q_2 and then using that $B < 1$, $q_1 > q_2$, and $2A - (B + 1) > 0$.

References

- Acemoglu, Daron, Simon Johnson, and James Robinson (2001) "The Colonial Origins of Comparative Development: An Empirical Investigation," *The American Economic Review*, Vol 91, No. 5, December, 1369-1401.
- Alesina, Alberto and Eliana La Ferrara (2002), "'Who Trusts Others?'," *Journal of Public Economics*, 85(2), 207-34.
- Alesina, Alberto and Paola Giuliano (2007), "The Power of the Family", Harvard University, mimeo
- Banfield, Edward C. (1958). *The Moral Basis of a Backward Society*. Free Press.
- Berg, Joyce, John Dickhaut and Kevin McCabe (1995), "Trust, Reciprocity and Social History", *Games and Economic Behavior*, 10: 122-42.
- Bisin, Alberto, Giorgio Topa and Thierry Verdier (2004). "Cooperation as a Transmitted Cultural Trait." New York University, mimeo.
- 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.
- Bisin, Alberto and Thierry Verdier (2001). "The Economics of Cultural Transmission and the Evolution of Preferences." *Journal of Economic Theory*. 97(1), 298-319.
- Cipriani, Marco, Paola Giuliano and Olivier Jeanne (2007). "Like Mother Like Son? Experimental Evidence on the Transmission of Values from Parents to Children." *IZA*. Discussion paper No. 2768
- Djankov, Simeon, Edward Glaeser, Rafael La Porta, and Florencio Lopez-de-Silanes and Andrei Shleifer (2003), "The New Comparative Economics", *Journal of Comparative Economics*, Vol. 31, 4, pp. 595-619
- Dohmen, Thomas, Armin Falk, David Huffman and Uwe Sunde (2007). "The Intergenerational Transmission of Risk and Trust Attitudes." *IZA*. Discussion Paper No. 2380
- Fehr, Ernst, Urs Fischbacher, Bernhard von Rosenblatt, Juergen Schupp, Gert Wagner, 2003, "A Nation-Wide Laboratory. Examining Trust and Trustworthiness by Integrating Behavioral Experiments into Representative Surveys", CESifo working paper 866.

- Fountain, Jane E. (1997) “ *Social capital: A Key Enabler of Innovation in Science and Technology*”, in L.M. Branscomb and J. Keller, eds., *Investing I Innovation: Toward a Consensus Strategy for Federal Technology Policy*, MIT Press Cambridge Mass.
- Glaeser, Edward, David Laibson, Josè A. Scheinkman and Christine L. Soutter, (2000), "Measuring Trust," *Quarterly Journal of Economics* 115(3): 811-846.
- Glaeser, Edward L., Bruce Sacerdote and Jose A. Scheinkman (1995), “Crime and Social Interactions”, *Quarterly Journal of Economics*, 111: 507-48
- Goldin, Claudia and Larry Katz (2001). ‘Human Capital and Social Capital: the Rise of Secondary Schooling in America, 1910–1940’, in (R. Rotberg, ed.), *Patterns of Social Capital*, Cambridge: Cambridge University Press.
- Guiso, Luigi and Monica Paiella (2007). “Risk Aversion, Wealth and Background Risk” *Journal of the European Economic Association*, forthcoming.
- Guiso, Luigi, Paola Sapienza and Luigi Zingales (2003). “People's Opium? Religion and Economic Attitudes.” *Journal of Monetary Economics*. 50(1), 225-82.
- Guiso, Luigi, Paola Sapienza and Luigi Zingales (2004). “The Role of Social Capital in Financial Development.” *The American Economic Review*. 94(3), 526-556.
- Guiso, Luigi, Paola Sapienza and Luigi Zingales (2005). “Cultural Biases in Economic Exchange.” *National Bureau of Economic Research*. Working paper, 11005.
- Guiso, Luigi, Paola Sapienza and Luigi Zingales (2006). “Does Culture Affect Economic Outcomes?”, *Journal of Economic Perspectives* (2006) 20: 23-48.
- Guiso, Luigi, Paola Sapienza and Luigi Zingales (2007). “*Long-Term Persistence*.” University of Chicago, mimeo.
- Guiso, Luigi, Paola Sapienza and Luigi Zingales (2008) , “Trusting the Stock Market” *Journal of Finance*, forthcoming.
- Knack, S. and Keefer, P. (1997). "Does social capital have an economic impact? A cross-country investigation", *Quarterly Journal of Economics*, vol. 112 (4), pp. 1252–88.2.
- La Porta, Rafael, Florencio Lopez de Silanes, Andrei Shleifer, and Robert Vishny (1997). "Trust in large organizations", *American Economic Review*, 87:2, 333-338..

- Putnam, Robert D. (1993). *Making Democracy Work*. Princeton University Press.
- Rice, Tom W. and Jan Feldman (1997). "Civic Culture and Democracy from Europe to America." *Journal of Politics*, 59, 1143-1172.
- Ruben Ernesto, Paola Sapienza, and Luigi Zingales (2007), Does Mistrust Breeds Mistrust? , University of Chicago working paper
- Sapienza, Paola, Anna Toldra, and Zingales Luigi (2007), "Understanding Trust," NBER working paper.
- Schupp, Jürgen, and Gert G. Wagner (2002). *Maintenance of and Innovation in Long-Term Panel Studies. The Case of the German Socio-Economic Panel (GSOEP)*. Allgemeines.
- Soroka, Stuart N., John F. Helliwell, and Richard Johnston. "Modeling and Measuring Trust." In *Diversity, Social Capital, and the Welfare State*, edited by Fiona Kay and Richard Johnston. In press, University of British Columbia Press.
- Tabellini, Guido (2005). "Culture and Institutions: Economic Development in the Regions of Europe." *IGIER*. Mimeo, Bocconi University.
- Tabellini, Guido (2007). "The Scope of cooperation: norms and incentives." *IGIER*. Mimeo, Bocconi University.
- Tabellini, Guido (2008). "Institutions and Culture", Presidential Address at the meetings of European Economic Association Meetings, Budapest August 2007 .
- Uslaner Eric M. (2007). *Where You Stand Depends Upon Where Your Grandparents Sat: The Inheritability of Generalized Trust*. Mimeo, University of Maryland – College Park.
- Williamson, Oliver E. (2000), "The New Institutional Economics: Taking Stock, Looking Ahead", *Journal of Economic Literature* Vol. 38, pp. 595-613..

Table 1: Trust and age (WVS data)

The table shows the relation between generalized trust and age using three waves of the World Values Survey (1981-4, 1990-3 and 1995-7). Generalized trust is based on the answers 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? The variable is equal to 1 if participants report that most people can be trusted and zero otherwise. “Health status” is coded based on the question: “All in all, how would you describe your state of health these days? (1=Very poor; 2=Poor, 3=Fair, 4=Good, 5=Very good)”. “Male” is an indicator variable equal to one if the respondent is male, otherwise equal to zero. “Age” is expressed in years. “Education” is the age in years at which the respondent completed his or her highest education (excluding apprenticeships). “Social status” is coded based on the response to the question: “People sometimes describe themselves as belonging to the working class, the middle class, or the upper or lower class. Would you describe yourself as belonging to the: 1=Lower class, 2=Working class, 3=Lower middle class, 4=Upper middle class, 5=Upper class”. “Income” is coded based on the response to the question: “Here is a scale of incomes. We would like to know in what group your household is, counting all wages, salaries, pensions and other incomes that come in. Just give the letter of the group your household falls into, before taxes and other deductions” (income categories are coded by decile for each society, 1=lowest decile, 10=highest decile). The religion variables are defined in detail in GSZ (2003) Table 1. Numbers in brackets are standard errors. The regressions include a country fixed effect and survey-year dummies.*** significant at less than 1% confidence level; ** significant at 5%; * significant at 10% level.

	Total sample	low trust countries (<25 th pct)	high trust countries (>25 th pct)
Log(age)	0.0495*** (0.0040)	0.0344*** (0.0066)	0.0556*** (0.0049)
Health status	0.0425*** (0.0018)	0.0232*** (0.0030)	0.0485*** (0.0021)
Male (0,1)	0.0022 (0.0030)	0.0026 (0.0049)	0.0027 (0.0037)
Education(years)	0.0057*** (0.0004)	0.0033*** (0.0005)	0.0072*** (0.0005)
Social status	0.0132*** (0.0015)	0.0043 (0.0030)	0.0149*** (0.0018)
Income decile	0.0082*** (0.0007)	0.0070*** (0.0013)	0.0083*** (0.0008)
Atheist	0.0315*** (0.0046)	0.0454*** (0.0110)	0.0308*** (0.0052)
Raised religiously	0.0163*** (0.0057)	0.0226** (0.0098)	0.0157** (0.0068)
Go to church at least once a year	0.0181*** (0.0059)	-0.0053 (0.0120)	0.0233*** (0.0068)
Go to church at least once a month	0.0321*** (0.0082)	0.0224 (0.0155)	0.0362*** (0.0096)
Observations	89677	22546	67131
R-squared	0.095	0.037	0.073
<i>t</i> -test for equality of age coefficient			2.585

Table 2: The trust-age profile at the country level

In column 1 the left hand side variable is the difference in the average generalized trust of the “old” (those older than 30) and the average generalized trust of the “young” (not older than 30); in column 2 it is the coefficient on age in an individual level regression of the trust of the individual on a number of characteristics with each regressions run separately for each country in the WVS. “Trust gap” is the difference between and index of average trustworthiness and the average trust of the young in the country. The trustworthiness index is obtained by averaging out the answers in the WVS about the individuals’ view on each of five types of behavior. Specifically individuals are asked “Tell me for each of the following statements whether you think it can be justified” (1= always, 10=never)”. We selected the following five statements to obtain a measure of average trustworthiness: “*Claiming government benefits to which you are not entitled*”; “*Avoiding a fare on public transport*”; “*Cheating on taxes if you have a chance*”; “*Buying something you knew was stolen*”; “*Accepting a bribe in the course of their duties*”. We have averaged out the answers to these questions and divided the result by 10 in order to have the variable on a fraction scale. We have then subtracted the share of young individuals in the country that say that generally speaking individuals can be trusted. *** significant at less than 1% confidence level; ** significant at 5%; * significant at 10% level.

	Trust old - trust young	Slope of the trust-age profile in country level regression
Trust gap	0.2267*** (0.0641)	0.0074*** (0.0022)
Trust gap* Mean trust of young	0.3679** (0.1784)	0.0220*** (0.0062)
Observations	53	53
R-squared	0.209	0.212

Table 3: The transmission of trust beliefs

In columns (1) (2) and (3) the left hand side variable is a dummy equal to 1 if the kid answers “Totally agree” or “agree slightly” to the question “What is your opinion on the following statement: On the whole one can trust people”; in columns (4) (5) and (6) it is a dummy equal to 1 if the kid answers “slightly disagree” or “totally disagree” to the question: “What is your opinion on the following statement: Nowadays one can't rely on anyone” ; in columns (7) (8) and (9) it is a dummy equal to 1 if the kid answers “slightly disagree” or “totally disagree” to the question: “What is your opinion on the following statement: If one is dealing with strangers, it is better to be careful before one can trust them”. Columns (3), (6) and (9) include as controls the age of the two paersnt and that of the son/daughter, their school attainment, a gender dummy and the income of the parents family and the son/daughter . *** significant at less than 1% confidence level; ** significant at 5%; * significant at 10% level.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Generalized trust	Generalized trust	Generalized trust	Reliance on others	Reliance on others	Reliance on others	Careful with others	Careful with others	Careful with others
Gen. trust_m	0.4668*** (0.0493)	0.4270*** (0.0537)	0.4128*** (0.0542)						
Gen. trust_f	0.3399*** (0.0489)	0.2530*** (0.0535)	0.2503*** (0.0538)						
Reliance_m				0.3527*** (0.0481)	0.2859*** (0.0508)	0.2587*** (0.0514)			
Reliance_f				0.3550*** (0.0481)	0.2969*** (0.0513)	0.2760*** (0.0518)			
Careful_m							0.6317*** (0.0760)	0.5772*** (0.0779)	0.5833*** (0.0786)
Careful_f							0.4084*** (0.0786)	0.3500*** (0.0806)	0.3235*** (0.0822)
Trustworthiness_m		0.0788 (0.0516)	0.0728 (0.0519)		0.2090*** (0.0504)	0.1917*** (0.0507)		0.1395** (0.0626)	0.1214* (0.0636)
Trustworthiness_f		0.1812*** (0.0517)	0.1690*** (0.0521)		0.1028** (0.0509)	0.0866* (0.0514)		0.2039*** (0.0621)	0.1753*** (0.0631)
Constant	-0.2418*** (0.0411)	-0.2885*** (0.0433)	-0.6400*** (0.1975)	-0.0763** (0.0372)	- (0.0409)	- (0.1986)	- (0.0329)	- (0.0506)	- (0.2552)
Additional controls	NO	NO	YES	NO	NO	YES	NO	NO	YES
Observations	3346	3279	3279	3344	3277	3277	3337	3271	3271

Figure 1. Timeline of decisions for generation t

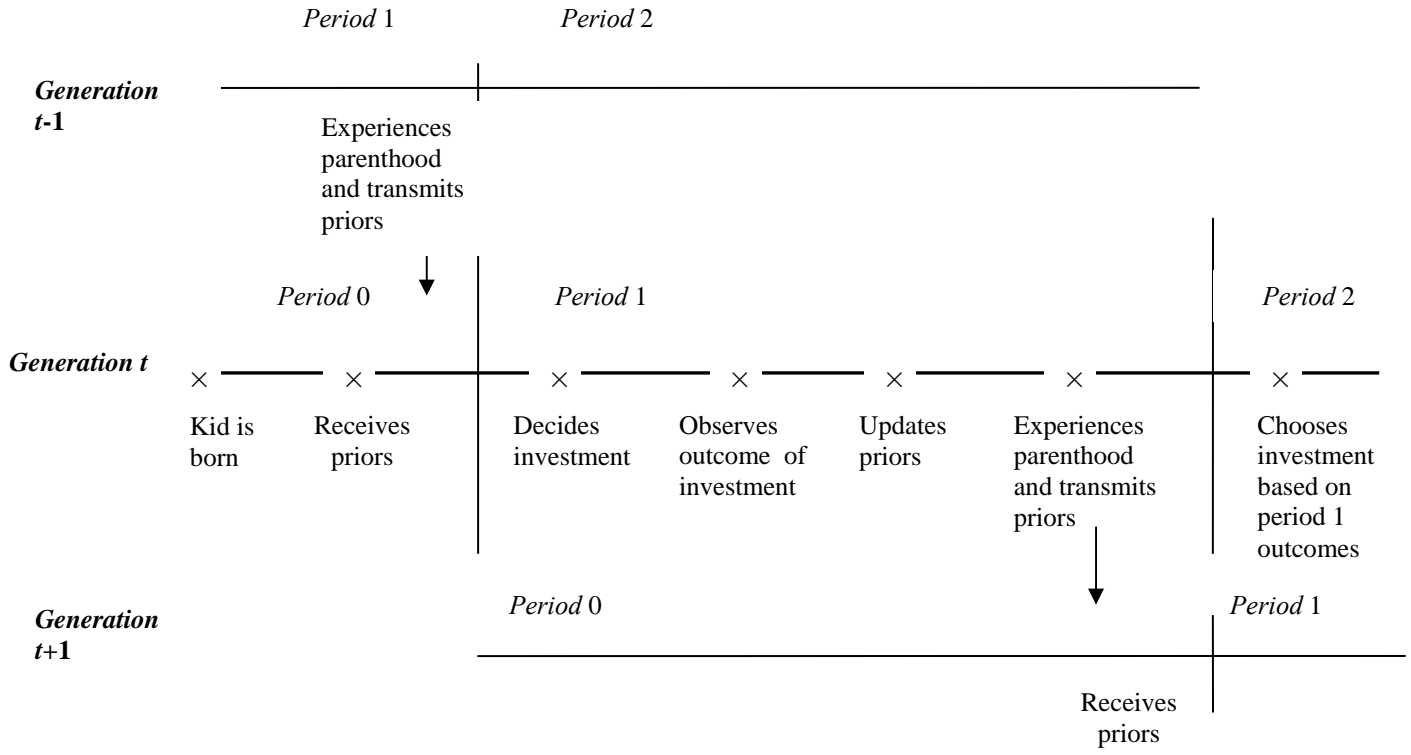


Figure 2. The transmission of priors.




Figure 3. Simulating the effects on average priors of a large shock to the return on trust.
The figure shows how the average prior evolves across generations following a shock to the return to cooperation.

Figure 3a:




Figure 3b:




Figure 3c:

