Theory Dec. (2013) 74:311–333
DOI 10.1007/s11238-011-9287-y

Trust, inequality and the market

Shaun P. Hargreaves Heap · Jonathan H. W. Tan · Daniel John Zizzo

Published online: 11 December 2011

Abstract This article examines, experimentally, whether inequality affects the social capital of trust in non-market and market settings. We consider three experimental treatments, one with equality, one with inequality but no knowledge of the income of other agents, and one with inequality and knowledge. Inequality, particularly when it is known, has a corrosive effect on trusting behaviours in this experiment. Agents appear to be less sensitive to known relative income differentials in markets than they are in the non-market settings, but trust in markets appears generally more vulnerable to the introduction of inequality than in the non-market setting.

Keywords Trust · Income inequality · Market · Social capital

JEL Classification C72 · C91 · Z13

1 Introduction

In liberal democratic societies, there is always a potential tension between the desire that people are treated equally in legal and political matters and the economic and social inequalities which arise through the exercise of individual freedom. Whether this tension is felt more or less acutely depends, in part, on what consequences flow, or are perceived to flow, from these inequalities. For this reason, and many others, there is a longstanding interest in economics on the impacts of economic inequality. This article is concerned with one potential mechanism through which inequality may affect

S. P. Hargreaves Heap · D. J. Zizzo (✉)
University of East Anglia, Norwich, UK
e-mail: d.zizzo@uea.ac.uk

J. H. W. Tan
University of Nottingham, Nottingham, UK

economic outcomes. It examines experimentally whether, and if so how, inequality affects the social capital of trust in two separate experimental settings: a market and a non-market one.

Evidence from survey data suggests that increasing inequality lowers trust and this, in turn, adversely affects economic and social outcomes. For example, Alesina and La Ferrara (2002) find that income inequality (and other types of heterogeneity among the relevant population) is associated with lower levels of trust, as revealed in survey data. There is also evidence that trust is a kind of social capital, in the sense that growth and a variety of other desirable outcomes improve when the level of trust found in country-level survey data rises (e.g. see Putnam et al. 1993; Knack and Keefer 1997; Beugelsdijk et al. 2004; Macours 2011).

While the links between inequality, trust and economic outcomes appear significant in these ways, the experimental evidence is less clear. First, there is laboratory evidence suggesting that what people say in a survey is not the same as what they do. Glaeser et al. (2000), for instance, find that reported attitudes to trust appear unrelated to whether subjects trust each other when playing a standard Berg et al. (1995) trust game (i.e., how much they give), although it is associated with their trustworthiness (i.e., how much they return). Second, while Hargreaves Heap and Zizzo (2009) find in an experiment that group heterogeneity reduces trust overall, Anderson et al. (2006), when focusing explicitly on heterogeneity due to inequality in a trust experiment, report little support for the hypothesis that inequality conspicuously lowers trust. Taken at face value, these experimental studies fail to cohere well with the stark conclusion that comes from the survey data linking inequality with low trust.

There are also reasons, however, for caution in drawing this negative inference from the experimental work; and these explain the design of our experiment. Typically, the survey data on trust is generated from answers to a general question: 'Generally speaking would you say that most people can be trusted or that you can't be too careful in dealing with people?' In this way, the survey is likely to test for a general inclination towards trust rather than one of trust in a particular person. The subjects of Glaeser et al. (2000), however, knew who they were playing the trust game with, and they identified common friends and social connections before playing; and so, arguably, their results do not directly address the more generalized sense of trust picked up in survey data.

The Anderson et al. (2006)'s subjects did interact anonymously, but the experiment only allowed for a weak form of knowledge of income inequality. Each subject knew only his or her income and the distribution of incomes, but not the income of the person they were interacting with. While this is sometimes the way that inequality is experienced, people frequently have a reasonably good idea of whether the person they are interacting with is richer, poorer or has about the same income. This follows from the way that role, dress, accent and the like can signal one's income even when the parties know little else about each other. There is also reason to suppose that knowing the other person's income might affect behaviour from a recent study by Lei and Vesely (2010) on intergroup biases between rich and poor groups and whether they survive a transition to a state of equality.

Our experiment addresses these issues. Our subjects play a trust game repeatedly and anonymously but with different co-players each time, and they do so both under

the condition of knowing only the distribution of incomes, and in a different treatment where subjects also know exactly the income of each person they are interacting with. The two treatments are compared against a baseline where incomes are equally distributed at a level equal to the average income in the inequality treatments. We thereby hope (a) to reproduce settings of sufficient anonymity where subjects are more likely to reveal the trust of the general kind that seems to be picked up by the survey data, (b) to match better the range of conditions under which inequality is often experienced, and (c) to compare directly populations with inequality versus those with equality using a between-subject design.

The other major contributions of the article come from the contrast between behaviours in the two institutional settings. One setting is supplied by the standard Berg et al. (1995) trust game, and the other is a version of the Fehr et al. (1993) experimental labour market game. This contrast is important in our design for two reasons.

First, there can be concerns over whether behaviour in a trust game in a laboratory is a good indication of how people actually behave in non-laboratory trust settings, particularly markets. One way of addressing this worry is to triangulate laboratory insights with those derived from field studies/experiments. Another is to examine experimentally how sensitive behaviour in the laboratory is to changes in the experimental setting in which the same kind of decision problem is embedded, particularly when those changes proxy what might be found in moving outside the laboratory. The two approaches are not mutually exclusive and we use the market game experiment in support of the latter strategy for testing robustness.

Second, the contrast between settings enables us to engage with the debate in economics over the influence of the institution of the market on people's behaviour. This literature has a long history too. Hirschman (1982) recounts, for example, an early set of arguments by Montesquieu, Condorcet, Hume and, in a more complicated way, Smith with respect to how the market was a powerful moralizing force in social relations when compared with more feudal attitudes which prevailed before the rise of the market. These are set against arguments like those of Marx which contend that markets create a set of false relations between people. In more recent times, the extension of the market relations, for example in childcare and through the introduction of (quasi) markets within the public sector, has rekindled this debate over whether markets are conducive or undermine the moral values that are their essential underpinning (see Folbre and Nelson 2000, and Lane 1991). The contrast between the two settings in our experiment allows us to examine a specific aspect of this debate: whether the institutional setting of the market affects the influence that inequality has on the social capital of trust.

Our experimental results support the conclusion drawn from the survey data: inequality does undermine the social capital of trust. The broad conclusion linking inequality with lower social capital is robust in the sense that, when people know about the level of inequality, it holds in both the market and the non-market settings. However, the fine-grained ways in which inequality affects individual behaviours varies across the two institutional settings. This suggests that generalizing the fine-grain details from the simple trust game to other settings may not be warranted. In particular, in the non-market version of the experiment trustworthiness, but not trust, decreases with known inequality, and the relative status (in terms of income) of the two players seems to

affect how much trust and trustworthiness there is between them. Conversely, relative status matters less in the market version of the experiment, but inequality appears to have a general corrosive effect in the sense that all income pairings tend to exhibit less trust and trustworthiness when there is inequality that is known. The verdict on the influence of the market is, therefore, mixed: people appear to be less sensitive to known relative income (i.e., status) differences in markets than they are in the non-market settings, but trust, and not only trustworthiness, is generally lower as the result of income inequality, even when the other agents' individual income is not known.

Section 2 describes the experiment. Section 3 sets out the results and Section 4 discusses them. Section 5 concludes.

2 The experiment

2.1 Outline

The experiment was conducted between May 2007 and 2008 at two universities.[1] Apart from the experimental instructions, the experiment was fully computerized. There were three treatments and each had 12 sessions, divided equally between the two universities. Each session had 12 subjects, and so a total of 432 subjects participated in the experiment. Subjects were randomly seated in the laboratory. Computer terminals were partitioned to avoid communication by facial or verbal means. Subjects read the experimental instructions and answered a control questionnaire, to check understanding of the instructions, before proceeding with the tasks. Experimental supervisors individually advised subjects with incorrect answers in the questionnaires. The experiment used 'experimental points' as currency, each worth either 1.9 UK pence (0.04 pounds) or 2.4 Euro cents (0.024 Euros) depending on location.[2]

The experiment was divided in two *stages*, each of which contained a number of *rounds*. In half of the sessions, subjects played trust games in the first stage and made (labour) market decisions in the second; in the other half, the reverse order applied. Counterbalancing operated in all treatments and at both locations to control for order effects.

Subjects were evenly assigned one of two roles at the beginning of the experiment, and they knew that everyone would keep their role throughout the experiment. 'J participants' were trusters in the trust games and the corresponding role (the employers) in the labour market game. 'K participants' were trustees in the trust games and the corresponding role (the workers) in the labour markets.

[1] The experiment was first run in Germany and then in England. Experimental instructions are provided in the electronic appendix A, which is available at http://www.uea.ac.uk/~ec601/TIM_EAppendices.pdf together with the other electronic appendices.

[2] The exchange rate of points into pounds was decided on the basis of the exchange rate between euro and pound at the time we began running the experiment in England.

2.2 The trust game stage

Subjects played six trust games, each time with a different co-player, and they knew that this was the case. These were standard Berg et al. (1995) trust games. The truster received an income in points (described below) and had to decide how many points (restricted to be between 0 and 10) to give to the other person and how many (if any) to keep. All the points given were multiplied by a factor of 3 before they were received by the trustee. The trustee then decided how much (if any) to keep and how much (if any) to return to the truster, who was informed of the outcome of the trustee's decision. The experimental instructions had a neutral frame (e.g., did not refer to 'trust', 'trusters', or 'trustees').

2.3 The (labour) market stage

We modeled our experimental market closely on Fehr et al.'s (1993) original paper on labour markets, and we followed them in choosing a neutral product market frame rather than using potentially loaded labour market labels in the instructions.[3] All subjects, six employers and six workers, made choices in ten rounds of a modified posted offer market that operated as follows.

They all received an income in points (described below) at the beginning of each round. Employers made wage offers between 0 and 40 points and could revise them at any point in the trading period as long a contract was not sealed. Workers had to choose whether to accept one or more of these wage offers. In a deviation from the Fehr et al. design, employers then approved (or not) one of the accepted contract offers. This mirrors real world markets where employers advertise a job ('wage offer'), workers apply for the job ('contract acceptance') but then employers choose which worker to hire for the job applied for ('contract approval'). Approval by employers sealed the contract. As in Fehr et al., employers could only employ one worker in any given round. Subjects had 4 min to seal contracts, with all live wage offers being displayed on all computer screens during this trading period. Examples of the computer displays shown in the experiment are provided in Fig. 1.

Once a contract was sealed, the worker had to decide on an effort level (between 0.1 and 1) at a cost in terms of experimental points which was ten times the effort level chosen. Therefore, if employed, the worker earned (on top of their initial income) the wage minus the effort cost in each round. Higher effort implied greater productivity and therefore bigger profits to employers. We adopted Fehr et al.'s formula to determine profits: in our case, this meant that profit equalled (40-wage) $\times$ (effort), which was added to their initial income.[4] Employers and workers who were not willing or able to seal a labour contract in any given round were left with their initial income.

[3] That meant using terms such as 'J participants', 'K participants', 'prices' and 'quality value' in place of 'employers', 'workers', 'wages' and 'effort'.

[4] Fehr et al.'s formula normalizes profits so as to avoid the possibility of net losses for the employer. This is especially useful in our case because, as we describe below, in two treatments employers have only a small endowment, and therefore there would be the danger of negative gains from the experiment resulting in possible distortions.

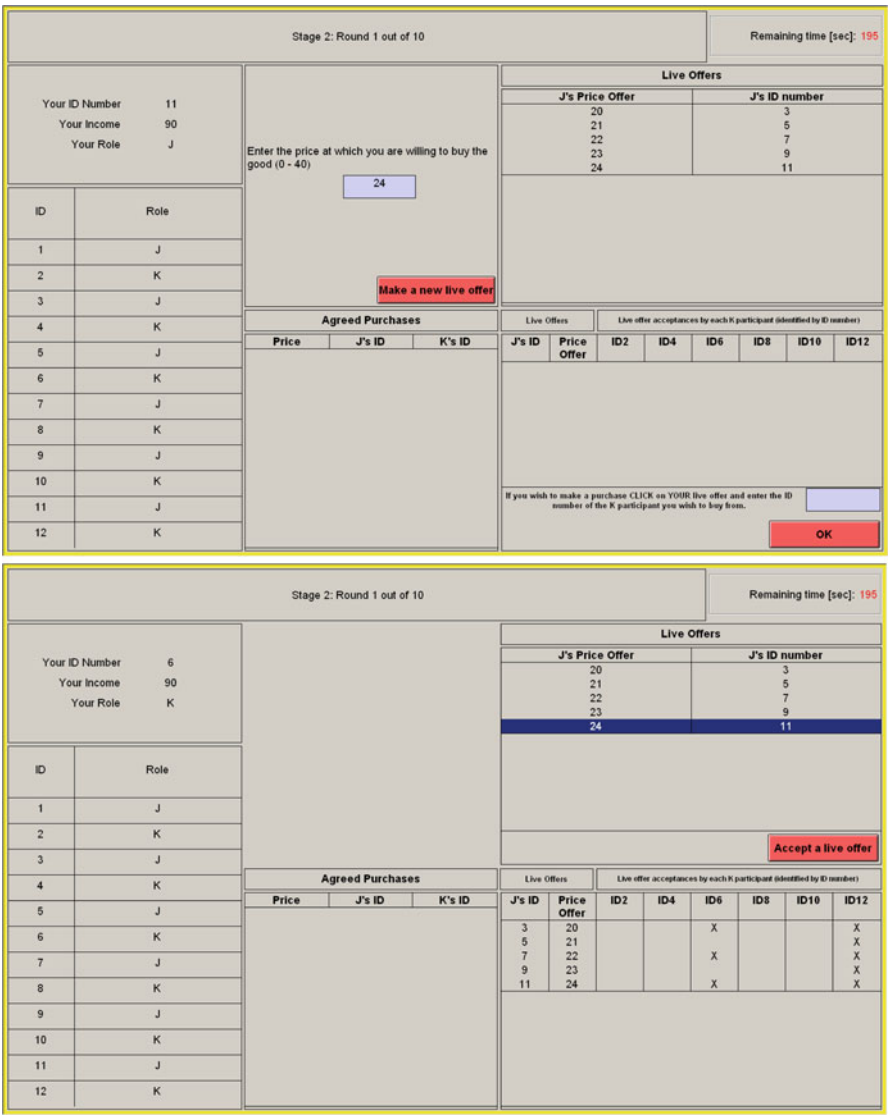

Fig. 1 Sample screens from the labour market stage

It is worth noting that, although the labour market game involves trust, the interaction is not isomorphic with that of the trust game. Subjects pair though choice in one and not the other and the agent whose decision affects the productivity of the interaction differs (it is the truster in the trust game whereas it is the equivalent of the returner, the employee, in the labour market game). These differences are for the most part not important as we will be comparing behaviour within either the trust or the labour market game as inequality changes and not across the two. However, they are potentially important because we do make some cross game comparisons; and in

this context it matters for interpretation that these differences arise largely because of the change to the market setting. This is the case with the first because it is a characteristic of markets that people are free to make contracts and so pair through choice. However, it would have been possible to construct a market exchange where the purchaser's decision affects productivity (through some effect via economies of scale on the conditions of production). This seems rather contrived and unusual as people in markets typically know the price when they decide on a purchase (and it is the seller's decision over quality that affects the productivity of the interaction rather than the buyer's). Fehr et al.'s (1993) game preserves this seemingly more natural property of markets and for this reason we use it. We return to this issue of interpretation when it arises in the discussion of the results.

2.4 Experimental treatments

There were three experimental treatments.

Treatment 1 is our baseline treatment with no inequality. All subjects received the same initial income of 50 points each round throughout the experiment, and knew this from the beginning.

Treatment 2 is our treatment with inequality in income but lack of knowledge about the specific level of income of other agents. At the beginning of the experiment, together with being assigned a role as a J or K participant, subjects were also given, in equal numbers, one of three possible levels of income: 10, 50 and 90 points per round (i.e., there were two subjects with each income level in each role). Each subject knew their own income as a result of this process and they also knew (a) the distribution of income across all subjects in the experiment and (b) that income, once assigned, would remain the same throughout the experiment. They did not know, however, the specific income level of other subjects that they interacted with: that is, whether they were 'rich' (with an income of 90), 'poor' (with an income of 10), or in the middle (with an income of 50). The differences in initial income levels were perspicuous (a) in the sense of being constant throughout the experiment (a rich subject knew he or she would remain rich throughout the experiment, and similarly for a poor or medium income subject) and (b) in the sense that, given the constraint on 'giving' of between 0 and 10 (or of offering a wage between 0 and 40), the richer player, when at his or her most generous, could at best only achieve equality of outcomes with the person in the next endowment bracket (e.g., in the trust game a subject with an income of 90 could give 10 to the trustee and, if he or she had 50 and returned nothing, each would get 50 points). These features combine to aid interpretability and external validity as real world wealth differences have considerable stability in time and individual transactions typically involve small transfers of wealth relative to the wealth differences between economic agents (e.g., it is not common for a rich person to give enough to a poor person so as to invert the wealth ranking). These features differentiate Treatment 2 from Anderson et al.'s (2006) experimental setup which it otherwise resembles.

Treatment 3 is our treatment with inequality in income and knowledge of the specific level of income of other agents. In the trust games, the computer displayed the income level of the coplayer. In the labour markets, the computer displayed the income level

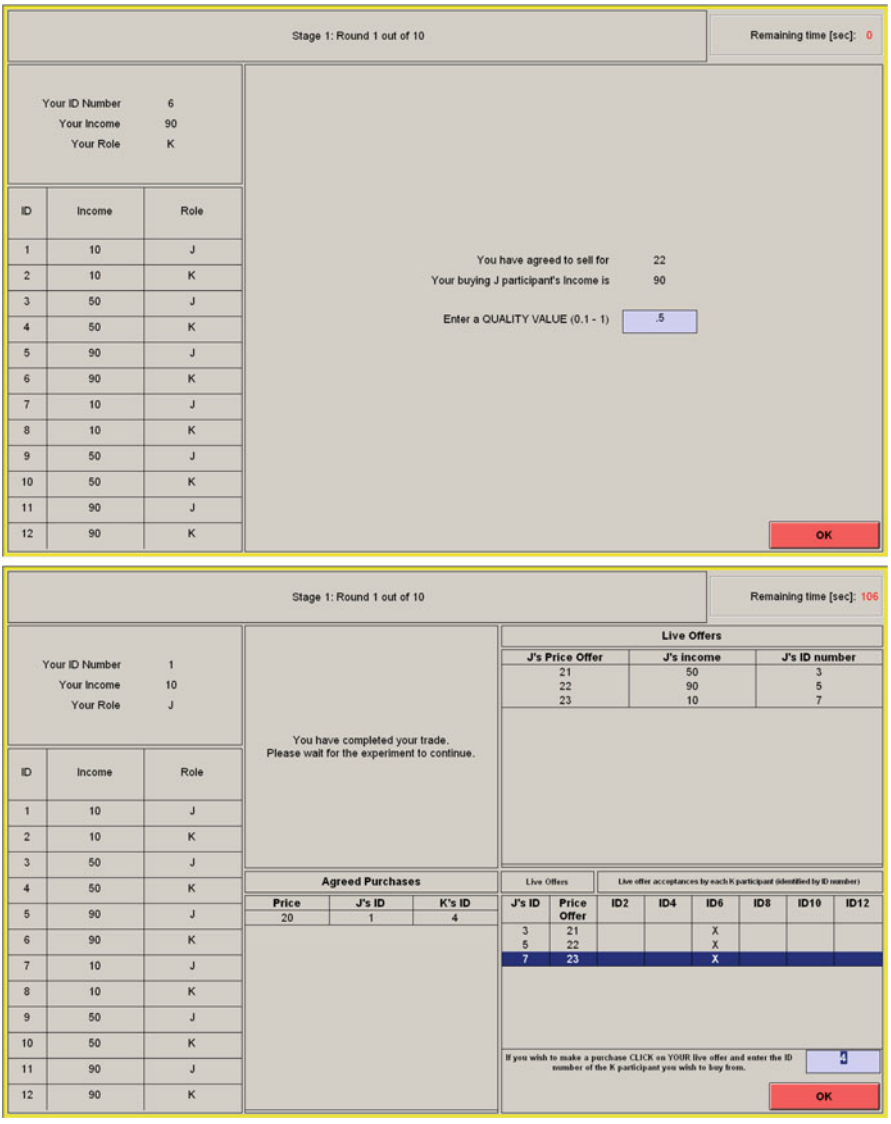

Fig. 2 Sample screens from treatment 3

of each employer making an offer and of each worker accepting an offer for everyone to see (Fig. 2 gives an example).

Taken together, the comparison between the three treatments enables us to test for whether inequality has a corrosive effect on trust and whether this changes when there is specific knowledge of inequality between agents both in terms of whether any such effect 'grows' or 'shrinks' and whether the manner of this effect changes with this knowledge. The point of the last observation about 'manner' is that inequality could affect all income pairings in the same way and/or it could depend for its effect on the

specific relative incomes (status) of each player in an interaction. Treatments 2 and 3 both allow for the operation of the first of these but only Treatment 3 gives scope for the second.

2.5 Predictions

In both the trust and market games, assuming self-interest and rationality, backward induction yields the unique equilibrium of minimal giving and wage offered by player J, and minimum return and effort exerted by player K, respectively. Deviations in behaviour by player J (K) thus measures trust (trustworthiness) in an economic sense.

We are not concerned in this paper with the motivational origins of trusting behaviours (we wish only to examine whether the social capital of trust is affected by inequality and not why), but there is one social preference that could be responsible for trust and which would produce changes in behaviour with inequality: inequality aversion (e.g., Fehr and Schmidt 1999). In this case, when there is inequality, trusting behaviours would be modified so as to reduce the inequality.

2.6 Payments

Each session lasted a little over one hour on average and the mean payment per subject was 17–18 pounds. Payments were based on the money earned and cumulated throughout the experiment. Subjects were privately paid and left the laboratory one at a time in an order designed to minimize the likelihood of seeing each other.

3 Experimental results

We consider the trust and labour market games in that order. In both, we begin by reporting descriptive statistics and univariate statistical tests on the differences in behaviour across treatments. This is followed by the results of regressions on individual behaviour that control for individual level and session level random effects, as well as for the effects of time.[5] Finally, we report on whether individual behaviour in the trust games is connected with that in the labour market game.

3.1 Behaviour in trust games

Figure 3 shows the giving rate, defined as the fraction of the endowment given by trusters to trustees, across treatments. This is the index of 'trust' in this game and using the Mann Whitney U test, taking the average giving rate in each session as an

[5] This specification controls for the possible non-independence of observations both at the individual level and at the session level, and is an efficient estimation method. Alternative regressions, such as those controlling for non-independence by robust clustering by sessions rather than random effects, yield very similar results. Therefore, none of our key results depend on the estimation method.

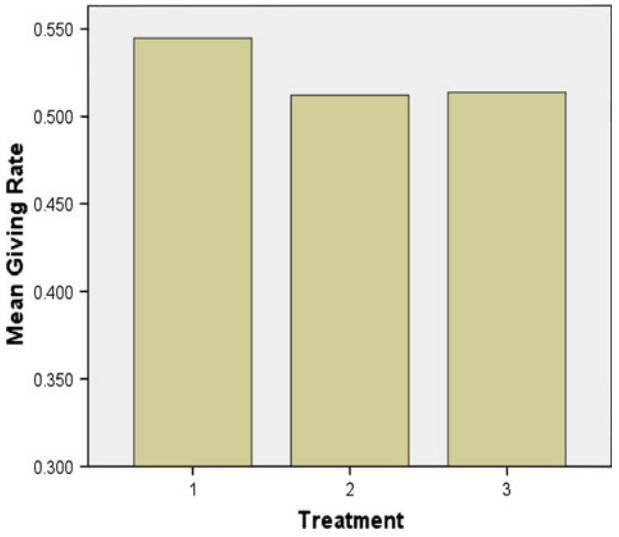

Fig. 3 Average giving rate across treatments

Table 1 Distribution of giving and return rates across pairings and treatments

	Treatment 1	Treatment 2	Treatment 3		
			To 10	To 50	To 90
Giving rate					
From 10		0.496	0.504	0.565	0.573
From 50	0.545	0.526	0.454	0.55	0.496
From 90		0.514	0.45	0.506	0.525
Return rate					
From 10		0.361	0.284	0.225	0.206
From 50	0.297	0.36	0.286	0.275	0.21
From 90		0.27	0.297	0.236	0.183

independent treatment, we find no statistically significant difference in giving rates across treatments.[6]

Result 1 Neither the introduction of inequality nor changes in the knowledge of this inequality has an effect on the average giving rate.

Table 1 decomposes the giving rate by the income endowment of the truster and in treatment 3 that of the trustee. The latter reveals a tendency in most cases of the truster to give more to a trustee who has a higher income and less to a trustee with a lower

[6] We also tested, here and elsewhere in the analysis, for location and counterbalancing effects, but were unable to find systematic patterns worth reporting. Men tend to have a higher giving rate.

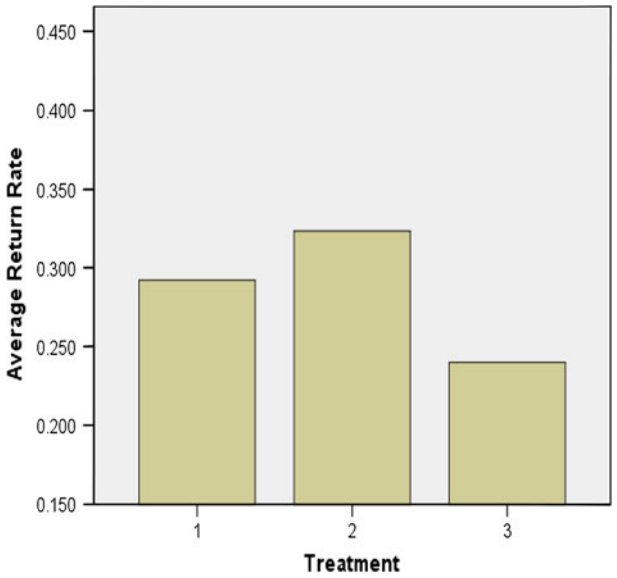

Fig. 4 Average return rate across treatments

income than to a trustee with same income and we examine this more closely in the regression analysis below.

Figure 4 reports on the return rate, defined as the fraction of the amount received from trusters that is returned by trustees (to the trusters), across treatments. This is the measure of 'trustworthiness' in this game. It is significantly higher (Mann Whitney U test, $P < 0.05$) in treatment 1 (where there is no inequality) than in treatment 3 (where there is inequality and specific knowledge of income level of coplayers), taking the average return rate in each session as an independent treatment.[7] The return rate is also significantly higher ($P < 0.01$) in treatment 2 (where there is inequality in income but lack of knowledge about the specific level of income of other agents) than in treatment 3 (where there is this knowledge of inequality). Overall, return rates go down by around 25% as we move from treatment 1 to treatment 3.

Result 2 The introduction of inequality that is known lowers the average return rate.

Table 1 also decomposes the return rates across the various income endowments of the trustees and in treatment 3 also by the income endowment of the truster. It shows that in treatment 2 there is some evidence that return rates decrease with the own income of the truster. This effect is absent in treatment 3 and instead we find evidence, as one reads across for the influence of the income of the truster on return rates, of inequality aversion in that return rates decrease as the truster's income rises. This effect holds across trustees of different incomes. By comparing the return rates across treatments, it seems that this inequality aversion effect operates in a negative

[7] All Mann Whitney tests below similarly control for the non-independence of observations within sessions by treating the session as the independent unit of observation.

direction of reducing the trustworthiness towards the rich, rather than increasing trustworthiness towards the poor (which seems no different to that observed in treatment 2). It is also not obvious that this effect is explained by reciprocity (i.e., by the standard relationship whereby trustees who are trusted more fulfill trust more),[8] since there is no evidence in the data that rich people trust less in the trust game (and so return less on grounds of reciprocation).

We now turn to the individual regression results in Table 2 that control for individual level and (nested in) session level random effects. They enable us to determine the effect of inequality on individual behaviour while controlling more systematically for a range of variables. The dependent variable for Models 1 and 1′ is the giving rate. The independent variables for Model 1 are *Location* (=1 for Germany, and=0 for UK), *Gender* (=1 if male, and=0 if female), and *EconBus* (=1 studying for an Economics and/or Business degree), *Round* (from 1 to 6), *NoInfo* (=1 for treatment 2, with inequality without information about co-player's income), *Info* (=1 for treatment 3, with inequality with information about co-player's income), *Income* (=50 in treatment 1 and =10, 50, 90 in treatments 2 and 3) and *DIOO* × *Info*, a variable that captures the difference between one's income and that of the co-player when this is known (i.e. interacted with *Info*). The *Info* and *Noinfo* dummies enable us to test for an effect from the existence of inequality under different information conditions in *general*: that is, how inequality affects all income pairings in the same way. Against this background, the *Income* and *DIOO* × *Info* variables allow us to examine whether the influence of inequality on behaviour depends, respectively, on one's own income and income relative to that of one's coplayer (i.e. it captures considerations of relative status that are induced by inequality). Model 1′ is the same except *NoInfo* is removed while *Info* remains so as to isolate the pure effect of having inequality and information against the benchmark provided by the other two treatments which do not have this combination.

The dependent variable for Models 2 and 2′ is the return rate. Model 2 has the same set of independent variables as Model 1, but also controls for reciprocity: it also has *Give*, which is how much, as a giving rate, the trustee has received before his or her decision to return, and *GiveSquared*, which is *Give* squared, to capture the possible non-linearity of this relationship.[9] Model 2′ has the same set of independent variables as Model 1′, but again adding *Give* and *GiveSquared*.

Round is negative and statistically significant in all four regressions, pointing to a decrease in giving and returning rates with time. Model 1 shows that neither *Info* nor *NoInfo* are significant, suggesting that there is no *general* effect on giving that comes from inequality and/or its knowledge (in the sense that all income pairings are affected in the same way). Nevertheless, model 1 does indicate that relative income (or status) matters for *individual* behaviour when there is inequality that is known: *DIOO* × *Info* (the relative income variable) is negative and significant ($P < 0.05$) meaning that the rich give less to the poor and the poor give more to the rich than they do, respec-

[8] This might occur for a number of psychological motives which have been documented in other experiments, such as inequality aversion (Fehr and Schmidt 1999), reciprocity (Falk and Fischbacher 2001) or trust responsiveness (Guerra and Zizzo 2004).

[9] We thank seminar participants in Jena for pointing us out this potential non-linearity.

Table 2 Regression analysis of giving and return rates in the trust game

	Giving rate					
	Model 1			Model 1′		
	β	t	P	β	t	P
Location	0.02	0.52	0.602	0.021	0.54	0.59
Gender	0.083	2.55	0.011	0.085	2.61	0.009
EconBus	−0.011	−0.3	0.764	−0.011	−0.31	0.753
Round	−0.01	−3.73	0	−0.01	−3.73	0
NoInfo	−0.019	−0.41	0.682			
Info	−0.03	−0.65	0.517	−0.021	−0.51	0.608
Income	0.0002	0.35	0.723	0.0002	0.32	0.748
DiooxInfo	−0.0008	−2.29	0.022	−0.0008	−2.28	0.023
Constant	0.574	9.62	0	0.564	10.22	0
	Return rate					
	Model 2			Model 2′		
	β	t	P	β	t	P
Give	0.144	1.6	0.11	0.148	1.62	0.105
GiveSquared	−0.062	−0.84	0.401	−0.066	−0.88	0.377
Location	0.016	0.66	0.506	0.02	0.79	0.429
Gender	0.019	0.84	0.404	0.02	0.81	0.415
EconBus	−0.032	−1.35	0.176	−0.033	−1.24	0.216
Round	−0.011	−5.71	0	−0.01	−4.73	0
NoInfo	0.037	1.3	0.195			
Info	−0.045	−1.68	0.093	−0.063	−2.52	0.012
Income	−0.0011	−2.46	0.014	−0.0012	−2.59	0.01
DiooxInfo	0.0012	4.56	0	0.0012	4.63	0
Constant	0.385	8.51	0	0.404	8.83	0

$n = 1296$; the models were estimated using GLLAMM (Rabe-Hesketh et al. 2004) and control for both session level and subject level random effects. DiooxInfo is the income differential between an agent and her or his co-player when this income differential is known, and as a result can take one of the following values: −80, −40, 0, 40 and 80 (when the income differential is unknown, it is equal to 0)

tively, to members of their own income group.[10] Other things being equal, the coefficient implies that low income subjects are given around 6–7% less than high-income subjects. This form of 'seignorage' when there is inequality nets out in the aggregate because we have symmetric distribution of incomes, and so does not produce a

[10] One should note that the magnitude of the marginal effect should be considered alongside the fact that the dependent variables of giving, returning and effort are rates constrained to the upper bound of 1, while independent variables such as Income and DIOOxInfo spans across 10–90 and −80 to 80, respectively.

difference in the average giving rates across treatments. Thus, this 'seignorage' effect is consistent with Fig. 3 and Result 1, and what is revealed in Table 1.[11]

Model 2′s coefficients on the giving rate or the squared giving rate are not as statistically significant as one might expect given the common finding of reciprocity in trust games (though *Give* does achieve significant at $P < 0.06$ one tailed). However, removing either *Give* variable makes the remaining one significant ($P < 0.001$) and so the lack of significance when both are present is probably an artifact of the correlation between the two variables. The lower *average* return rate when there is known inequality (Result 2) appears to be explained in these regressions by two effects that inequality has on individual behaviour in treatment 3 and one in treatment 2. First, the return rates decrease with *Income* ($P = 0.015$). Thus, the rich return less when there is inequality and this contributes to a lower average return rate in *both* treatments 2 and 3. Second, when inequality is known there is a general corrosive effect of inequality on trustworthiness, with an additional 6% drop in the return rate in treatment 3 relative to the two others (model 2′, coefficient on *Info* is negative and significant, $P = 0.012$). With both effects operating in treatment 3, but only the former in treatment 2, it is not surprising to find that the average giving rate is lower in treatment 3.

In addition, and in contrast with giving behaviour, return rates increase with the difference in income (*DIOO* × *Info* is positive and significant, $P < 0.001$). This relative status effect is perhaps more familiar than what was found in giving behaviours: it corresponds to a form of inequality aversion, or, to use a feudal description, there is a form of 'noblesse oblige' whereas there was 'seignorage' in giving. The influence of 'noblesse oblige' on the *average* return rate, like that of 'seignorage' when there is a symmetric distribution of income, nets out to produce no overall effect. The size of the coefficient on the relative income term is similar to the negative one on *Income* and so when combined with the influence of relative income, the rich return an especially small amount to each other and the poor are not returned more than the average in other treatments (and this is what is suggested in Table 1). These effects together imply that overall inequality, when known, can only decrease and never increase overall trustworthiness.

In these ways the individual regression results both support the findings reported in Results 1 and 2 with respect to how the presence of inequality affects average giving and return rates. They reveal some specific status or relative income related effects on individual behaviour.

Result 3 The introduction of inequality has a general corrosive effect on trustworthiness and there are individual level relative status or income effects both in 'giving' (in the form of 'seignorage') and in 'returning' (where there is a form of inequality aversion).

[11] The only additional significant variable is gender in these equations: it seems that men give more than women.

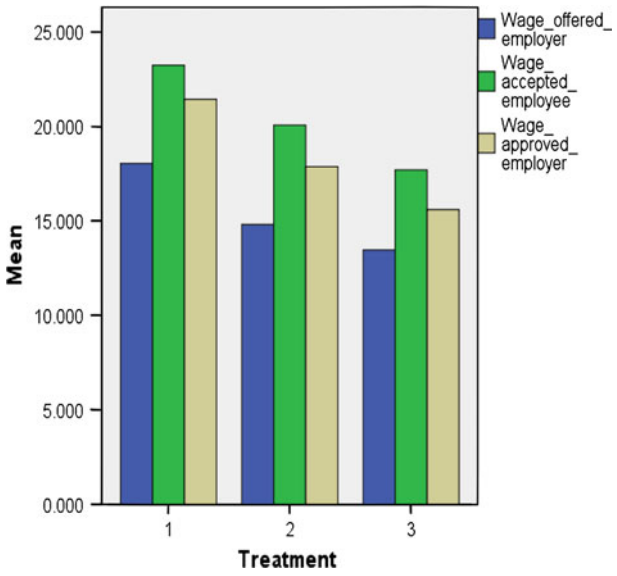

Fig. 5 Average wage offered, accepted and approved across treatments

3.2 Behaviour in the labour market game

'Trust' in the labour market game is captured by the size of the wage offered and approved by the firm once a contract has been accepted by the worker. Figure 5 shows that the *wages* offered, accepted and approved decrease from treatment 1 (equality) to 2 (inequality without knowledge of the coplayer's wage) and to 3 (inequality with knowledge about the co-player's wage). The decrease in wage approved between treatment 1 and 3 is about 25% and the difference across treatments is statistically significant ($P < 0.03$ in a Kruskal Wallis test). Table 3 performs the disaggregation of this treatment data by the income of the firm and, in treatment 3, the worker. There are no consistent patterns in the wage approved in relation to income of the firm or relative income.

Result 4 The introduction of inequality lowers the average wage.

Figure 6 reports the number of offered, accepted and approved contracts. Inequality has no effect here on the likelihood of employment/unemployment as the number of approved contracts remains close to 6. However, there is a decline in the number of contracts offered (and accepted) when there is inequality in treatments 2 and 3, respectively, relative to treatment 1 (Mann Whitney $P < 0.05$ and $P < 0.1$, respectively). In this sense, there is less market search when there is inequality.

'Trustworthiness' in this game is measured by the effort of workers and this is reported in Fig. 7. There is no statistical difference between treatments 1 and 2, but there is significantly lower effort, by around 20–25%, in treatment 3 relative to either treatment 1 or 2 (Mann Whitney $P = 0.02$).

Table 3 Distribution of approved wages and work efforts across pairings and treatments

	Treatment 1	Treatment 2	Treatment 3		
			To 10	To 50	To 90
Wage approved					
From 10		15.876	13.55	14.1	13.991
From 50	21.419	17.956	16.817	16.357	15.729
From 90		17.396	16.092	15.508	15.93
Effort					
From 10		0.42	0.319	0.331	0.286
From 50	0.389	0.29	0.299	0.287	0.261
From 90		0.391	0.232	0.276	0.315

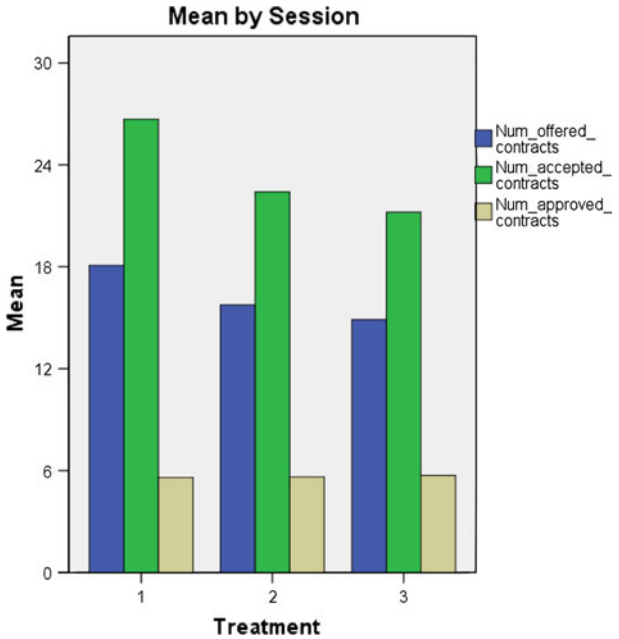

Fig. 6 Number of contracts offered, accepted and approved across treatments

Result 5 Average effort is lower when there is inequality that is known.

The distribution of average effort exerted for each type of pairing (in terms of relative income) is given in Table 3. While effort is generally lower in treatment 3 than in the other two, there is no pattern which suggests inequality aversion. The effort level by low income workers is constant regardless of employer income. Middle income employees exert about the same effort when there is inequality, regardless of whether or not the income of the employer is known, when compared with case without inequality. High-income employees exert more effort when paired with a high-income employer as compared with a pairing with a low income employer (Mann Whitney $P = 0.02$).

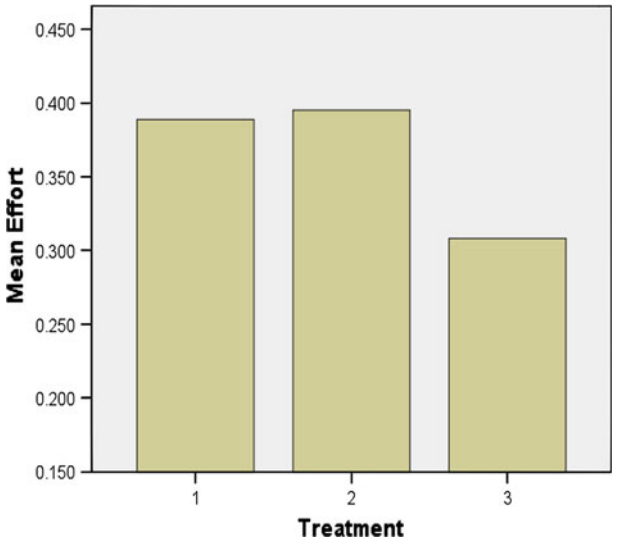

Fig. 7 Average effort exerted across treatments

Table 4 gives the regression results on individual behaviour, controlling for individual level and (nested in) session level random effects. The dependent variable for Models 3 and 3′ is the approved wage rate, and for Models 4 and 4′ the effort level. The independent variables are the same as those in the trust game regressions, with the exception that Models 4 and 4′ control for *Wage*, the wage rate received by the employee when he or she makes his or her effort decision, and *WageSquared* to control for non-linearities.[12]

Info and *NoInfo* are both negative and statistically significant (both $P < 0.001$) in models 3 and 3′ suggesting that the introduction of inequality, whether with or without knowledge of the coplayer's income, has a general corrosive effect on the level of trust revealed in the market part of this experiment. This in turn largely explains at the individual level why the average wage level is lower when there is inequality (as in Fig. 6; Result 4). There is one offsetting effect from inequality that is picked up in 3 and 3'. The *Income* variable is positive and marginally significant ($P < 0.06$). This would imply that the rich employers offer better wages. There are two other results which also contrast with those found in the trust game. First, we find that *DIOO* × *Info* (the relative income variable) is *not* statistically significant. Second, *Round* is now positive and statistically significant, pointing to a general increase in individual wages over time ($P < 0.001$).[13]

In Models 4 and 4′, effort tends to become lower as the experiments get closer to the end (Round is significant and negative, $P < 0.001$). There is marginal support

[12] *Round*, of course, now ranges from 1 to 10.

[13] We also find a small educational background effect, insofar as employers with an Economics and/or Business background offer lower wages ($P < 0.05$), and a small location effect, as wages were slightly lower in the U.K. ($P < 0.001$).

Table 4 Regression analysis of wages and effort in the labour market game

	Wage					
	Model 3			Model 3′		
	β	t	P	β	t	P
Location	−1.923	−3.49	0	−2.055	−3.68	0
Gender	−0.124	−0.24	0.808	0.025	0.05	0.96
EconBus	−1.33	−2.2	0.028	−1.26	−2.46	0.014
Round	0.342	9.29	0	0.333	8.53	0
NoInfo	−3.25	−4.52	0			
Info	−4.303	−7.33	0	−4.14	−7.48	0
Income	0.0198	1.93	0.054	0.0242	2.54	0.011
DiooxInfo	0.0043	0.7	0.486	0.0034	0.55	0.581
Constant	17.261	22.09	0	16.939	19.41	0
	Effort					
	Model 4			Model 4′		
	β	t	P	β	t	P
Wage	0.01	4.05	0	0.01	4.07	0
WageSquared	0	1.51	0.132	0	1.52	0.129
Location	0.018	0.58	0.56	0.017	0.59	0.557
Gender	0.054	1.77	0.078	0.053	1.89	0.059
EconBus	0.016	0.44	0.66	0.015	0.46	0.648
Round	−0.008	−5.69	0	−0.008	−5.74	0
NoInfo	0.002	0.06	0.952			
Info	−0.053	−1.14	0.255	−0.055	−1.61	0.108
Income	−0.0011	−2.32	0.021	−0.0011	−2.4	0.017
DiooxInfo	0.0005	1.94	0.053	0.0005	1.97	0.049
Constant	0.244	3.97	0	0.247	4.9	0

$n = 2033$; the models were estimated using GLLAMM (Rabe-Hesketh et al. 2004) and control for both session level and subject level random effects. DiooxInfo is the income differential between an agent and her or his co-player when this income differential is known, and as a result can take one of the following values: −80, −40, 0, 40 and 80 (when the income differential is unknown, it is equal to 0)

for a general corrosive effect from inequality (*Info* is negative but weakly significant, $P < 0.06$, one tailed). We find that in both models, however, effort increases significantly with *Wage*, the wage received ($P < 0.001$), which is a standard result (e.g. Camerer 2003). It is this finding that probably, largely explains Result 5: that is, this reciprocating behaviour when combined with the lower average wage offers when there is inequality (Result 4) explains the lower average effort level under inequality. Rich workers make less effort (*Income* is negative and significant, $P < 0.05$). There is some evidence of relative income effect (*DIOO* × *Info* is positive and significant, $P < 0.06$) although this is not as large in magnitude as in Models 2 and 2′.

Result 6 Lower wages and effort under inequality largely arise from the general corrosive effect that the presence of inequality has on all individual wage offers and the influence of reciprocation in individual effort decisions. Relative income effects appear to matter less in the market setting than in the non-market one.

We conclude by testing the prediction that behaviour in the trust game can be used to predict behaviour in the labour market game. Mean Spearman ρ by subject between the giving rate and wage offered or approved, or the number of approved contracts are around 0.11–0.13 ($P < 0.1$ or better).[14] The mean ρ between return rate and effort is high at 0.338 ($P < 0.001$). Regression analysis controlling for individual level and session level random effects suggests that giving rates and wage offers are not correlated, while, again, trustworthiness in the trust game is a good positive predictor of market effort.

Result 7 Trustworthiness in the trust game is a good predictor of trustworthiness in the market game.

4 Discussion

These results are important in several respects. First, inequality, particularly when it is known, does appear to undermine the social capital of trust. In the trust game, subjects were less trustworthy on average when there was inequality that was known (Result 2). Both trust and trustworthiness are lower in the labour market game when there is known inequality (Results 4, 5). This conclusion lends support to the findings from the survey evidence on the relation between trust and inequality. It stands in contrast to the earlier experimental work of Anderson et al. (2006) which was, at best, weakly consistent with the survey evidence.

We conjectured that we might find a difference of this sort in our experiment because we allowed in one treatment for the subjects to have specific knowledge of each other's income endowments; whereas the earlier experimental work by Anderson et al. had only allowed a general knowledge of the inequality within the whole experimental group. The evidence of our experiment is consistent with this conjecture because we find that there are less trusting behaviours when subjects know the income endowments of their coplayers. This result, that knowledge of the coplayer's income has a negative effect, is again important because it would help justify, for example, privacy policies that keep knowledge of individual incomes in companies where trust might matter a confidential matter. It also provides a rationale against policies that make income differentials salient.

[14] Specifically, mean ρ(giving rate, wage offered)=0.114 ($P < 0.1$); mean ρ(giving rate, wage approved)=0.120 ($P < 0.1$); mean ρ(giving rate, number of approved contracts)=0.128 ($P < 0.05$). Spearman correlation coefficients are computed for each subject and then averaged out across the session to find a mean Spearman correlation by session: P values are computed using nonparametric sign tests in relation to this mean correlation statistic. This procedure enables us to control for the non-independence of observations within each session.

Second, the implications of this experiment for the external validity, particularly in relation to market behaviours, of this (and other) trust game experiments are mixed. The fact that we find that inequality has a corrosive effect in both settings adds weight to the broad inference that the social capital of trust is not served by inequality. The fine-grained effects of inequality are, however, more complex across the settings. For example, while trustworthiness falls by a similar amount in each setting when there is inequality, this decrease is compounded in the market setting by a lower level of initial trust when there is inequality (see Results 4, 5 and 6 as compared with Results 1 and 2). This difference casts some doubt on whether the fine-grained details of the results concerning individual behaviour found in this and other simple trust game experiments can be generalized to other settings where trust decisions are made (e.g. those that occur in the market). That being said, while more research is clearly called for, we found it encouraging for research on trust games that trustworthiness in the trust game does seem to generalize to market settings (see Result 7).

Third, the institution of the market appears to have interesting effects on individual behaviour in the presence of inequality in this experiment. On the one hand, inequality is more corrosive of trust in the market setting. On the other, there is notably less evidence that people treat other people differently on the basis of their relative income or status differentials in our market version of the experiment, whereas in the non-market setting we find that individuals treat each other differently depending on their status and rank (see Result 3). There are traces of what, in terms of feudal attitudes, would be termed 'seignorage' and 'noblesse oblige' in the non-market game but not in the same degree in the market one. In this sense, the shift to the market context may make income inequality matter more in terms of a general corrosive effect on trust, but it also, apparently, encourages individuals to treat each other more equally. The latter lends direct support to what Hirschman (1982) identifies as the *doux commerce* thesis regarding the influence of the market (and this, too, contrasts with an earlier experimental study of the influence of status in market transactions—see Ball et al. 2001). In addition, and insofar as our result adds weight to more general arguments of this kind, it can be thought to lend support to the traditional argument that markets are likely to undermine differences in individual discriminatory behaviour (see Becker 1971). These inferences regarding the influence of the market are, of course, qualified by the earlier observation that there is a change between the two versions of the game that cannot be strictly associated with shift to the market (i.e. the agent whose decision affects productivity). We make them, nevertheless, because it is not obvious how this difference would by itself account for the differences in behaviour (i.e. how it would interact with inequality).

There is another respect in which it is significant that our experimental design allows us to examine whether inequality has a general effect on all interactions and/or effects that are specific to relative income or status of the players. This is because the interest that individuals might have in reducing inequality in order to promote trust is likely to be different in these two cases. Everyone stands to gain from a reduction in inequality when inequality has a general effect on all interactions between individuals. This is not the case when inequality produces relative income or status related behaviours. In such cases and if these relative income behaviours are weakened by reductions in inequality, then some will gain and some will lose (in the case of giving).

This difference is likely to be of some significance for the politics of inequality and it is not one that can be easily explored through the survey data.[15]

Finally, there is evidence in both institutional settings that individual trust and trustworthiness decisions are affected differently by inequality. The apparent existence of the different feudal attitudes that attach to each type of trusting decision in the non-market setting is a case in point. In addition, there is the difference in the way that trust is not affected, on average, by inequality in the non-market trust game but, when the income of the coplayer is known, trustworthiness is. The market reframes this differently, but an asymmetry between trust and trustworthiness remains: it is the trust decision that is affected by inequality and the effect on trustworthiness is a derivative of this. In short, our subjects do not seem to think about trust and trustworthiness in the same way when there is inequality which is known; and this is consistent with other experimental findings (see Glaeser et al. 2000).

5 Conclusion

In this article we report on an experiment that examines how the presence of inequality might affect the social capital of trust. This is potentially important because the social capital of trust has been found to be related to a range of desirable economic and social outcomes and, while there is considerable survey evidence on why trust might vary across populations, there is little experimental evidence. We find that inequality undermines trusting behaviours in our experiment. In this respect, the experiment supports the inference from the survey data that there is a link from inequality to economic performance that comes through the influence that inequality has on trust. Crucially, we also find that knowledge of the coplayer's income matters. This has obvious further policy implications insofar, for example, as it might justify statutory and human resources policies that keep income differentials confidential.

The experiment also addresses the methodological question that is concerned with the external validity of experimental results. This is assuming increasing importance among experimentalists and non-experimentalists alike (see Guala 2005). We do this through embedding similar trust decisions in two institutional settings in our experiment: one market and the other a non-market one. The broad result that inequality undermines the social capital of trust looks more robust than it otherwise would precisely because we find that this is the case in both institutional settings. Nevertheless, there are differences in the way that inequality affects individual behaviour in our two settings and this should caution against any attempt to generalize (to other settings including those outside the labouratory) the specific features of individual behaviour

[15] To be specific it would seem that markets create, in this respect at least, a stronger interest in egalitarianism than do non-markets. There is a further respect in which these experiments complement the findings from the survey data. It is possible, in principle, that the observed differences in average trust when there is inequality arise from the aggregation of what is actually an effect that income has on individual behaviour (i.e. it is not that inequality affects individual behaviour, it is income which when aggregated over different income distributions creates a relation between trust in the aggregate and income inequality). Alesina and La Ferrara (2002) control for this possibility in their analysis of the survey data and still find that inequality has an effect on trust. Our experiment does the same and comes to a similar conclusion.

that are revealed in this (and other) simple trust games. This, too, is important. In one respect, though, we find an ability to generalize from one setting to the other: namely, trustworthiness in the non-market trust game is a good predictor of trustworthiness in the market setting.

The final contribution of this article comes from the difference in the effect of inequality on individual behaviour in our two institutional settings. This is interesting for what it suggests about the influence of the market on individual behaviour. There is a long history to the debate over this influence and, perhaps characteristically, we find two rather opposing effects. Markets appear from our experiment to help encourage people to treat each other in the same way. That is, people in markets are less guided by their relative income and status than is the case in the non-market setting. Against this, however, people seem to be less willing to trust each other as the result of the mere existence of income inequality and this is the case even when the other agents' individual income is not known.

Acknowledgments We wish to thank Johannes Abeler, Simon Gächter, Martin Sefton, Pieter Serneels, Bob Sugden, Chris Starmer and participants to presentations in East Anglia, Jena, Lyon and Nottingham for useful advice, Marta Sernec, Eliza Tan, Kei Tsutsui, James Watson and Helen Gilbert for research assistance. Part of the experiments was conducted during Jonathan Tan's employment at the European University Viadrina. We thank the Institute of Microeconomics, European University Viadrina, and the University of East Anglia for financial support. The usual disclaimer applies. The experimental instructions are found at http://www.uea.ac.uk/~ec601/TIM_EAppendices.pdf.

References

Alesina, A., & La Ferrara, E. (2002). Who trusts others? *Journal of Public Economics, 85*, 207–234.

Anderson, L., Mellor, J., & Milyo, J. (2006). Induced heterogeneity in trust experiments. *Experimental Economics, 9*, 223–235.

Ball, S., Eckel, C., Grossman, P., & Zame, W. (2001). Status in markets. *Quarterly Journal of Economics, 16*, 161–188.

Becker, G. S. (1971). *The economics of discrimination*. Chicago: Chicago University Press.

Berg, J., Dickaut, J., & McCabe, K. (1995). Trust, reciprocity and social history. *Games and Economic Behaviour, 10*, 122–142.

Beugelsdijk, S., deGroot, H., & van Schaik, A. (2004). Trust and economic growth: Robustness analysis. *Oxford Economic Papers, 56*, 118–134.

Camerer, C. (2003). *Behavioural game theory*. Princeton: Princeton University Press.

Falk, A., & Fischbacher, U. (2001). Distributional consequences and intentions in a model of reciprocity. *Annales d'Economie Et de Statistique, 63-64*, 39–62.

Fehr, E., Kirchsteiger, G., & Riedl, A. (1993). Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics, 108*, 437–460.

Fehr, E., & Schmidt, K. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics, 114*, 817–868.

Folbre, N., & Nelson, J. (2000). For Love or money—Or both?. *Journal of Economic Perspectives, 14*, 123–140.

Glaeser, E., Laibson, D., Scheinkman, J., & Soutter, C. (2000). Measuring trust. *Quarterly Journal of Economics, 115*, 811–846.

Guala, F. (2005). *The methodology of experimental economics*. New York: Cambridge University Press.

Guerra, G., & Zizzo, D. J. (2004). Trust responsiveness and beliefs. *Journal of Economic Behaviour and Organization, 55*, 25–30.

Hargreaves Heap, S., & Zizzo, D. J. (2009). The value of groups. *American Economic Review, 99*, 295–323.

Hirsch, F. (1977). *The social limits to growth*. London: Routledge and Kegan Paul.

Hirschman, A. (1982). Rival interpretations of market society: Civilizing, destructive, or feeble?. *Journal of Economic Literature, 20*, 1463–1484.

Knack, S., & Keefer, P. (1997). Does social capital have and economic pay-off? Cross-country investigation. *Quarterly Journal of Economics, 112*, 1251–1288.

Lane, R. (1991). *The market experience*. Cambridge: Cambridge University Press.

Lei, V., & Vesely, F. (2010). In-group versus out-group trust: The impact of income inequality. *Southern Economic Journal, 76*, 1049–1063.

Macours, K. (2011). Increasing inequality and civil conflict in Nepal. *Oxford Economic Papers, 63*, 1–26.

Putnam, R., Leonardi, R., & Nanetti, R.Y. (1993). *Making democracy work: Civic traditions in modern Italy*. Princeton: Princeton University Press.

Rabe-Hesketh, S., Skrondal, A., & Pickles, A. (2004). GLLAMM Manual. U.C. Berkeley Division of Biostatistics Working Paper 160, Berkeley.