Negative Income Shocks Increase Discount Rates*

Johannes Haushofer†, Daniel Schunk‡, Ernst Fehr§

September 16, 2013

Abstract

Poor individuals often exhibit higher discount rates than richer individuals, and negative income shocks have been linked to increases in discounting. However, it remains unknown whether this effect is causal, and whether is due to beliefs or preferences: increases in discount rates in poverty or after negative income shocks could either be due to (beliefs about) more severe environmental constraints in contexts of poverty, or to direct effects of poverty or income shocks on preferences. Here we address these questions in a laboratory experiment in which subjects receive different starting endowments which create "rich" and "poor" groups. All participants then perform a real effort task to earn money, following which subgroups of participants receive positive and negative income shocks. Subsequently we measure discount rates, effort provision, social preferences, reservation wages, and psychological outcomes. We find that negative income shocks lead to an increase in time discounting. In contrast, positive income shocks weakly decrease discount rates. Discount rates are not affected by mean levels of wealth ("rich" vs. "poor"). The effect of negative income shocks on the discount rate is specific to short-run discounting, i.e. negative income shocks increase *decreasing* impatience. No strong effects were found on effort provision, social preferences, reservation wages, and psychological outcomes. Together, these findings suggest that income shocks have a direct causal effect on discount rates that is not explained by wealth levels.

*We thank Florian Bosshart, Amélie Brune, and Deborah Kistler for excellent research assistance. This research was supported by Cogito Foundation Grant R-116/10 and NIH Grant R01AG039297 to Johannes Haushofer.

†Program in Economics, History, and Politics, Harvard University, Abdul Latif Jameel Poverty Action Lab, MIT, and Department of Economics, University of Zurich. joha@mit.edu

‡Department of Economics, Mainz University. daniel.schunk@uni-mainz.de

§Department of Economics, University of Zurich. ernst.fehr@econ.uzh.ch

1 Introduction

Both empirically and theoretically, poverty has been linked to present-focused discount rates (see literature review below). However, it remains unknown to what extent this relationship is causal; in addition, it is unknown whether an effect of poverty on discount rates is due to beliefs about environmental constraints, or changes in preferences. Evidence exists for both views. In favor of the former hypothesis, it has frequently been argued that high rates of time discounting may be entirely rational in conditions of poverty – e.g., Becker & Mulligan (1997) show how existing economic conditions such as wealth, and environmental influences such as mortality and risk, can endogenously lead to behavior that looks like impatience. Others have argued that market imperfections, e.g. in insurance and credit markets, are responsible for behavior that resembles high rates of time discounting in poor countries, since they may imply that investments with immediate fixed costs and delayed benefits are less attractive than immediate-return investments (Banerjee, 2001; Holden et al., 1998; Pagiola, 1996, Shiferaw & Holden, 2001; Bernard et al., 2011). Similarly, present-focused decisions may be rational if returns to investment are low, which has also been argued to be the case in developing countries (Rosenzweig, 1995; van Walle, 2003). Present-focused behavior among the poor may also occur if people are severely calorie-constrained and close to subsistence (Murphree, 1993; Bardhan, 1996; Lumley, 1997; Dasgupta, 1997). Finally, Rachid (2012) has recently argued that the poor may show present-focued discount rates in order to avoid thinking about their typically gloomy future. Endogenous discount rates may have significant long-run consequences: Strulik (2011) argued that such endogenous discount rates can account for the growth paths taken by many countries better than the standard Ramsey model, and Chakrabarty (2011) suggested that they can create poverty traps. However, it could also be the case that conditions of poverty also affect discount rates directly, above and beyond a rational response to environmental conditions. A number of authors have suggested that environments of poverty may have direct impacts on preferences (Bertrand et al., 2004; Hall, 2008). Here we present evidence for this view; specifically, we show that negative income shocks – a salient feature of the lives of many of the world's poorest people – increase discount rates, and particularly decreasing impatience (Epper et al., 2001; Prelec, 2004). Because of the difficulty of studying income shocks while holding constant wealth, this paper takes a laboratory approach, which allows to study the effect of shocks while holding constant individual wealth by a) random assignment to treatment condition, and b) comparison of an income shock group to a control group with the same absolute wealth level.

Poverty and Discount Rates

Numerous papers document empirically that a relationship exists between discount rates and poverty. In one of the earliest studies of this kind, Hausman (1979) found that poorer households had higher discount rates than richer households. Lawrance (1991) estimated discount rates from Euler equations in a panel of US households, and oconfirmed this finding. Similar results were found by Sullivan (2011), who report weak evidence in a sample from China that households with lower wealth have higher discount rates. Pender & Walker (1990) elicited discount rates from Indian peasants, and found that they were significantly predicted by the wealth of the participants, again with wealthier subjects exhibiting less discounting. Similar results were found by Yesuf & Bluffstone (2008), who elicited discount rates from 262 households in rural Ethiopia and found that these were significantly negatively correlatd with wealth variables such as the amount of land owned, the value of a household's capital stock, and, to some degree, the number of oxen owned. Note, however, that a number of studies show no significant relationship between poverty and discount rates (Cagetti, 2003; Gourinchas & Parker, 2002; Stephens & Krupka, 2006; Ogaki & Atkeson, 1997).

The findings summarized above suggest that may be at least a correlational inverse relationship between discount rates and poverty. To test the causal role of poverty in affecting decision-making, Tanaka et al. (2010) elicited time preferences from 181 individuals in Vietnam and fit a number of discounting models to the data. They find that income is negatively related exponential discounting; interestingly and in contrast to our findings, no such relationship is evident for hyperbolic discounting, which loosely corresponds to the measure of decreasing impatience we use here (Epper et al., 2011; Prelec, 2004). A particularly elegant feature of this study is the use of rainfall as an instrumental variable for income: the negative relationship between exponential (but not hyperbolic) discounting and income persists when income is instrumented using the exogenous rainfall data, and thereby establishes a causal channel from low income to high discounting. A similar approach was taken by di Falco et al. (2011), who found in a sample of 1237 households in rural Ethiopia that income shocks caused by severe droughts led to increases in the rate of time preference. Note, however, that it is not possible in these settings to disentangle the income effect of the shock from any psychological effects of the shock itself; i.e., increased rates of time preference after a shock may be entirely rational responses to the changed economic circumstances, rather than a genuine change in preferences.

In a study related to ours, Spears (2010) randomly assigned poor participants in India to one of four conditions: they could either be rich or poor, in the sense that they received either two (rich) or one (poor) good from a choice set of three options. In addition, each participant could either be in a "choice" or a "no choice" condition, where the former meant

that participants could choose which item(s) from the choice set they wanted to receive, whereas in the latter case the items were randomly assigned. Spears then asked participants to perform two tasks that are frequently used to measure cognitive control: squeezing two handlebars for as long as possible, and performing a "Stroop"-type task. In this task, participants have to name the number of items in a display, which, confusingly, are numbers themselves; thus, a display might be "3 3", in which case the correct answer would be "two" since there are two items in the display. Spears found that participants performed worse on the two cognitive-control tasks if they had been randomly assigned to both the poor and the choice conditions. Thus, requiring people to make a choice appeared to deplete cognitive control, but this was only true for poor participants. Since cognitive control has frequently been related to hyperbolic discounting in the psychology literature (Shamosh et al., 2008; Shamosh & Gray, 2008), the study by Spears (2010) is similar in spirit to ours; however, note that while Spears focuses on low absolute levels of income, we hold levels constant and focus instead on negative income shocks. In addition, we measure economic choice rather than performance in psychological games.

Thus, a number of studies have suggested that conditions of poverty are associated with increased discount rates. Two questions arise: first, is this relationship causal? Second, is it due to beliefs about environmental constraints, or changes in preferences? A number of authors have suggested that environments of poverty may have direct impacts on preferences (Bertrand et al., 2004; Hall, 2008). Here we present evidence for this view; specifically, we show that negative income shocks – a salient feature of the lives of many of the world's poorest people – increase discount rates, holding constant absolute levels of wealth. Because of the difficulty of studying income shocks while holding wealth constants, this paper takes a laboratory approach, which allows to study the effect of shocks while controlling individual wealth by a) random assignment to treatment condition, and b) comparison of an income shock group to a control group with the same absolute wealth level.

Emotion and Discount Rates

Our paper is related to the literature that studies the effect of emotions on discount rates: if negative income shocks affect participants' affect, any effects on discount rates may be mediated by this channel. Indeed, emotions can have strong influences on discount rates; Loewenstein (1996, 2000) points out that in the presence of visceral factors such as rage, people sometimes exhibit extreme discounting of future events. Laboratory experiments that randomly assigned participants to particular emotion induction conditions confirm this hypothesis. For instance, Raeva et al. (2010) studied the effect of experienced regret on

discount rates. They induced regret in participants by first letting them choose one of two lotteries, and then revealing that the other choice would have produced a better outcome. Control participants only saw the result of their chosen lottery, not that of the alternative, and thus could not experience regret. After this manipulation, participants completed a discounting task; those participants who had experienced regret had a lower indifference point for an outcome that was available tomorrow, i.e. were less patient than control subjects. Conversely, participants who had experienced rejoicing, i.e. their chosen lottery was superior to the alternative, were more patient than controls. A similar paper by Ifcher & Zarghamee (AER, forthcoming) showed that positive affect, induced by a video clip of a stand-up comedian, made participants more patient than a control clip showing nature scenes.

Reference points and Discount Rates

In addition, this paper is also related to the literature on reference-dependent utility and prior outcomes. However, prior work makes conflicting predictions about the effect of income shocks on discount rates. On the one hand, existing evidence suggests that people evaluate options against existing reference points (Kőszegi & Rabin, 2006, 2007). In our experiment, it is likely that unexpected negative income shocks in our experiment put participants below the reference point (irrespective of whether that reference point is participants' own pre-shock income, or the average income of the group as a whole). Thus, the negative income shock should induce a loss frame, while unexpected positive income shocks should put participants above the reference point, i.e, in a gain frame. In prospect theory, making decisions in the loss domain is associated with risk-seeking behavior, while the gain domain is associated with risk aversion[1] (Kahneman & Tversky, 1979; Payne et al., 1980; Hershey & Schoemaker, 1980; Hershey et al., 1982; Slovic et al., 1982). Together with the theoretical and empirical relationship between risk and discount rates, according to which risk aversion induces present-focused discount rates (Leigh, 1986; Rachlin et al., 1991; Anderhub et al., 2001; Myerson et al., 2003; Eckel et al., 2004), this literature predicts that negative income shocks should make subjects risk-seeking and therefore lead to a decrease in the rate of time discounting, i.e. a preference for delayed outcomes.

On the other hand, the effect of prior losses on subsequent risky choice described by Thaler & Johnson (1990) makes the opposite prediction. These authors suggest that participants may "edit" the options available to them before making a choice; according to the hedonic editing rule, they do this in such a way as to make the resulting prospects appear most pleasant. In particular, this rule dictates that gains are segregated (i.e. considered

[1]Note that risk-seeking in the loss domain has been called into question by recent studies; e.g. Bruhin et al. (2010), Epper & Fehr-Duda (2012).

independently from prior outcomes), losses are integrated with prior outcomes. The implication of this rule is risk aversion after previous losses. Thaler & Johnson present empirical evidence in favor of this view: prior monetary losses lead to increased risk aversion among their participants. In the context of the present paper, this finding suggests that negative income shocks should in fact lead to more present-focused discount rates.

In sum, it remains open to what extent poverty and negative income shocks affect discount rates. We test here whether positive or negative income shocks, or absolute levels of wealth, affect discount rates. To this end, subjects receive different starting endowments which create "rich" and "poor" groups; in addition, all participants then perform a real effort task to earn money, following which subgroups of participants receive positive and negative income shocks. Subsequently we measure discount rates. We find that negative income shocks lead to an increase in discount rates, while positive income shocks lead to a weak decrease. Discount rates are not affected by persistent levels of wealth ("rich" vs. "poor"). The effect of negative income shocks on discount rates is specific to *decreasing* impatience (Prelec, 1989, 2004). Together, these findings suggest that negative income shocks may increase discount rates.

2 Methods

Participants

We recruited 148 healthy male participants form the subject pool of the University of Zürich. Their mean age was 22 ± 2.47 years (mean ± S.D.). We excluded students of economics and psychology. All participants gave written informed consent and received a show-up fee of CHF 10, in addition to any earnings from the experimental tasks, as described below. An experimental session lasted 2h. Participation was resctried to men because we also measured levels of stress hormones during the experiment, and controlling for ovarian cycle in women is logistically difficult. The neuroendocrine data are not reported here because no differences between experimental conditions were found. Participants were German native speakers. To ensure that they would be able to receive delayed payments, we included only participants who indicated that they would stay in Zurich at least for the subsequent 12 months.

Procedure

At the beginning of the experiment, each participant was informed about the nature of the tasks to be performed, as described below. After these instructions, each participant completed a PANAS questionnaire, which measures positive and negative affect (Watson et al.,

1988), and five visual analog scales, which asked to what extent participants currently felt a) stressed, b) in control of their lives, c) optimistic, d) self-confident, e) that the government should take responsibility for people's well-being, rather than individuals themselves. Participants marked their current feelings on a 10 cm line; responses were coded as between 0 and 100.

Each participant was randomly assigned to one of four treatment conditions, unbeknownst to them: "always rich"; "always poor"; "negative income shock"; "positive income shock". When the experiment began, participants in the "always rich" and "negative income shock" groups had a high initial endowment of 1000 points; in contrast, the "always poor" and "positive income shock" groups had a low initial endowment of 100 points. 70 points were converted into 1 CHF (USD 1.06 at the time of the study) at the end of the experiment and paid out.

Throughout the experiment, participants were informed of their own current wealth through bars and numbers on the screen; the size of the bar corresponded to the current wealth of the participant. In addition, bars were also shown for current maximum wealth, minimum wealth, and average wealth across all participants within the particular session. Thus, participants could continually keep track of their own wealth, and its relation to the wealth of the entire group of participants in their session. Bars were always normalized to the maximum wealth bar for ease of display.

Tasks

Real effort task Participants then participated in a real effort task for 15 periods, which resembled that used by Abeler et al. (2009 Abeler Felk Goette Huffman). Each period lasted 2 minutes. The task consisted of counting the number of zeros in a 7×5 random table to zeros and ones, which was presented on the left side of the screen. The right side of the screen displayed the wealth variables described above – own wealth, and maximum, minimum, and average wealth of all participants; the purpose of displaying this information even during task performance was to make own wealth in comparison to that of the entire group as salient as possible. After counting the zeros in a given table, participants entered their answer in a text field at the bottom of the screen. The next table was then displayed, without feedback about performance to minimize learning effects. Participants counted as many tables as they could within each 2 minute period, and earned 5 points for every correctly counted table. After each period, the accumulated points from the period were added to the wealth of the participant and displayed for 20 seconds in the middle of the screen, again also showing minimum, maximum, and average wealth. After these 20 seconds, the next period began.

Income shocks Participants played 15 periods of the real effort task, which lasted 35 minutes. After 15 periods of earning income, the two income shock groups received their income shocks. The timing, magnitude, and direction of these shocks was unanticipated; however, participants were informed at the beginning of the experiment that they might experience a sudden change in their wealth levels. Specifically, during the instruction period at the beginning of the study, participants were told that they might experience a change in their wealth during the real effort task that they would perform. They were told that they would experience either exactly zero or exactly one such wealth change, but were not told the timing, magnitude, or direction of this change. All participants were told that such sudden changes in wealth levels during the experiment were possible, even those in the "always rich" and "always poor" conditions; participants in these groups did not receive income shocks after period 15, nor were they told when the other participants experienced the income shocks. No justification was given for the income shocks; participants were informed of the shock through a screen that read "Your income has decreased by x points" or "Your income has increased by y points".

The magnitude and direction of the income shock for the "negative income shock" group was such that the post-shock average wealth of this group was equal to the pre-shock average wealth of the "always poor" group. Similarly, the magnitude and direction of the income shock for the "positive income shock" group was such that the post-shock average wealth of this group was equal to the pre-shock average wealth of the "always rich" group. Put differently, the two groups switched positions from the "poor" into the "rich" group, and vice-versa. This allows us to compare the effect of income shocks on economic choice, holding constant current wealth: comparing the behavior of the "negative income shock" group to the "always poor" group reveals the effect of a negative income shock, holding constant current wealth, while comparing the behavior of the "positive income shock" group to the "always rich" group reveals the effect of a positive income shock, again holding constant current wealth.

After receiving the income shock, participants were again presented with their updated wealth and the maximum, minimum, and average wealth across participants. This information was displayed for one minute to make their new wealth salient to participants in the shock groups. Participants then played two more periods of the real effort task; the purpose of these two periods was again to make participants fully aware of their new wealth situation and their position relative to others.

After period 17, participants performed the behavioral tasks of interest. The following sections describe these tasks in greater detail.

Intertemporal Choice Task Participants performed three blocks of an intertemporal choice task with varying delays, where decisions between a sooner smaller reward and a later larger reward were offered. In two of these blocks, subjects had the choice between a smaller reward tomorrow, and a larger reward in a) 6 months and 1 day, or b) 12 months and 1 day. The short delay was set to "tomorrow" rather than "today" to keep transaction costs the same for sooner and later payments. In the third block, subjects chose between a smaller reward in 6 months and 1 day, and a larger reward in 12 months and 1 day. Each block consisted of 6 binary choice trials, resulting in a total of 18 trials. The larger reward was kept constant at an amount of 30 Swiss Francs (CHF), while the sooner smaller reward started at CHF 15 and was then adjusted with a titration method according to the choices the subject made. Possible serial correlation and order effects in subjects' responses were controlled for by randomizing the order of trials across blocks, i.e. the order in which the various indifference points were determined. We presented subjects with choices in terms of CHF instead of points in this task to make the discounting task as distinct as possible from the effort task, in an effort to be conservative and minimize spillovers across tasks.

Reimbursement consisted of a flat rate of CHF 10 and a variable payment depending on participants' choices. In particular, as was explained to the participants at the beginning of the study, one of all their choices in the time preference task was randomly selected at the end of the study, and the chosen option on this trial was paid out, i.e., participants could pick up the chosen amount on the chosen day of delivery, using a voucher valid at the University cashier's office. Transaction costs were kept constant by setting the soonest outcome to "tomorrow".

Titration is a standard method for identifying discount rates in the discounting literature (Mazur, 1988; Green & Myerson, 2004; Kable & Glimcher, 2007; Rachlin et al., 1991). The titration used a bisection algorithm which set the initial small, soon amount for each delay combination to 50% of the large amount, and then gradually approximated the participant's indifference points for the different delay combinations[2]. The titration procedure lasted for 6 trials at each combination of delays; this means that each indifference point was identified to a precision of CHF 0.23 (CHF 15×0.5^6), i.e. the initial difference between CHF 15 and

[2]For each choice of the later reward, the sooner reward was increased by half the difference between it and 30 CHF; for instance, if a subject chose CHF 30 in 12 months and 1 day over CHF 15 tomorrow, the next trial would offer the subject a choice between CHF 30 in 12 months and 1 day and CHF 22.50 tomorrow; if the subject still chose CHF 30 in 12 months and 1 day, the next offer would be CHF 30 in 12 months and 1 day vs. CHF 26.25 tomorrow, and so on. For each choice of the sooner reward, the sooner reward was decreased by half of the difference between it and the previously offered soon reward. For instance, if a subject chose CHF 15 tomorrow over CHF 30 in 12 months and 1 day, the next trial would offer the subject a choice between CHF 7.50 tomorrow and CHF 30 in 12 months and 1 day; if the subject chose CHF 7.50 tomorrow, the next offer would be CHF 3.75 tomorrow vs. CHF 30 in 12 months and 1 day, and so on.

CHF 30/ CHF 0 was halved six times). The amount of the sooner reward at the end of this titration procedure was taken as the indifference point for the particular delay combination, i.e. the amount of the sooner smaller reward where participants switched between the smaller sooner and the later larger reward.

This procedure resulted in individual indifference points for each subject, which was then used as the basis for the computation of discount rates and indices for decreasing impatience and departures from stationarity, described below. Note that this procedure collapses subjects' choices in the discounting task into one or two parameters; thus, each subject entered the statistical analysis only once, i.e. we are not using multiple (non-independent) data points for each subject.

Calculation of discount rates For every subject and every delay level, we determined the amount at which a subject was indifferent between the earlier and the later option based on the individual indifference points (see above). This allowed us to express the subjective value of the delayed reward as a fraction of the subjective value of the immediate reward.

In our analysis, we consider the following measures of discounting, following Prelec (2004) and Epper et al. (2011). First, we use the indifference points themselves as outcome variables; these represent simply the subjective value of CHF 30 at a particular timepoint, as seen from an earlier timepoint. We then compute the exponential decay that is implied by each indifference point, separately for the 3 indifference points. For each participant i and each delay combination $(t_1 = 0, t_2 = 6)$, $(t_1 = 0, t_2 = 12)$, and $(t_1 = 6, t_2 = 12)$, and the amount of the larger, later payment, x_2:

$$x_1 = \exp\left(-\delta_{t_1,t_2}\frac{t_2 - t_1}{12}\right) x_2.$$

This implies:

$$\delta_{t_1,t_2} = -\frac{12}{t_2 - t_1}\ln\frac{x_1}{x_2}.$$

Next, we compute an index for decreasing impatience by taking the following difference:

$$\Delta_{DI} = \delta_{0,6} - \delta_{0,12}.$$

Further, we compute an index for departures from stationarity:

$$\Delta_{DS} = \delta_{0,6} - \delta_{6,12}.$$

The intuition for these indices is as follows: first, decreasing impatience implies less discounting over a long time horizon than a short time horizon; in our case, decreasing impatience

would predict a smaller annualized discount rate over a time horizon of 12 months compared to 6 months. The index Δ_{DI} captures this difference; it is larger when subjects discount more over a short time horizon (6 months) than over a longer time horizon (12 months). Similarly, a departure from stationarity implies less discounting over a period of 6 months when that period occurs in the distant future compared to when it occurs in the near future; in our case, a departure from stationarity would predict less discounting between 6 months and 12 months than between tomorrow and 6 months. The variable Δ_{DS} measures this difference, with larger values implying a greater difference between discounting in the near vs. the distant future, and hence a greater departure from stationarity.

Social Preference Task We further investigated whether and to what extent income shocks affected participants' social preferences; for instance, participants who had experienced a negative income shock might be less prosocial than others. We employed a modified version of the Ring Task (Buckley et al., 2001). In this task, participants decide between two options on each screen; each option contains a payoff to the participant, and a payoff to a randomly matched other participant. For instance, a particpant might decide between 600 points for themselves and 0 for the other participant, or 550 points for themselves and 230 for the other participant. Participants answered 12 such questions, which are summarized in Table 1. The social preferences of each participant were summarized by counting the number of prosocial decisions (0-12) on this task, where a prosocial decision was defined as one in which the participant chose a larger payoff for the other participant and the cost of a smaller payoff for themselves.

BDM Auction Task We next asked whether the income shocks changed participants' reservation wage; for instance, might participants who just lost a substantial proportion of their wealth be more willing to work at a lower wage? We therefore conducted a Becker-deGroot-Marschak (1964) auction in which participants could bid against the computer on the opportunity to complete the real effort task for another eight periods. Participants entered their bid into a text field, the computer played the auction immediately, and winning participants performed the real effort task for another 8 periods, while the remainder of the subjects waited until they had completed the experiment. The advantage of this type of auction is that it is incentive-compatible and elicits subjects' true willingness to pay for playing a further eight periods. The computer bid was randomly drawn from a uniform distribution between 0 and the expected earnings from a further 8 periods of play, based on performance of each subject in the first 15 periods of the real effort task. If participants' willingness exceeded the computer bid, they could perform the real effort task for the remaining 8 peri-

ods. At the end of the study, participants completed another PANAS questionnaire and five visual-analog scales (see above). Finally, they completed a socioeconomic questionnaire and the Barratt Impulsiveness Scale (Patton et al., 1995), and were paid and excused.

Econometric specifications

The effect of negative income shocks on the outcome variables was assessed using OLS regressions of the following form:

$$y_i = \beta_0 + \beta_1 \text{NEGATIVE SHOCK}_i + \beta_2 \text{POSITIVE SHOCK}_i + \beta_2 \text{ALWAYS RICH}_i + \gamma \mathbf{X_i} + \varepsilon_i \tag{1}$$

where y_i are the outcome variables decsribed above, NEGATIVE SHOCK$_i$, POSITIVE SHOCK$_i$, and ALWAYS RICH$_i$ are dummy variables indicating whether subject i was in the "negative income shock", the "positive income shock", or "always rich" group. The omitted category is the "always poor" condition. $\mathbf{X_i}$ is a vector of control variables which include yearly family income, a dummy for being currently in debt, and a dummy for being employed. ε_i is the error term.

3 Results

Effect of Income Shocks and Income Differences on Discount Rates

To ascertain that the income shock manipulations worked as intended, we first report the evolution of wealth levels while performing the real effort task. The "always rich" and "negative income shock" groups started the experiment with an endowment of 1000 points (CHF 14.28); during the first 15 periods, the average wealth level in these two groups grew to 1948.38 ± 28.60 (mean $\pm$ SEM) points, with no significant difference between groups (as is expected, since the groups were identical up to that point; always rich: 1923.78 ± 39.25; negative income shock: 1972.97 ± 41.75; $t = -0.86$, $p = 0.394$). Similarly, the "always poor" and "positive income shock" groups started the experiment with an endowment of 100 points (CHF 1.43); during the first 15 periods, the average wealth level in these two groups grew to 1029.46 ± 27.17 points, again with no significant difference between the groups (always poor: 1057.30 ± 44.97; positive income shock: 1001.62 ± 30.48; $t = -1.02$, $p = 0.309$). The magnitude and direction of the income shock was -918.92 ± 5.84 for the "negative income shock" group, and $+918.92 \pm 5.84$ for the "positive income shock" group. Note that these shocks are equal in magnitude and opposite in sign by design, since the two groups simply

switched positions; i.e., each participant in the "negative income shock" group lost the same number of points, and each participant in the "positive income shock" group gained the same number of points. The non-zero variance of the income shocks stems from the fact that the pre-shock difference between the groups differed somewhat across experimental sessions. In sum, the real effort task and the experimental manipulation of wealth levels through income shocks worked as intended. Figure [fig:inc_shock] shows the evolution of wealth levels as a function of period throughout the experiment; it can be seen that the post-shock wealth levels match exactly those of the "always rich" and "always poor" groups, respectively.

The main question of this study was whether income shocks affect discount rates, while wealth levels are held constant. Our design allows us to test this hypothesis as follows: first, comparing the "negative income shock" group to the "always poor" group after the income shock identifies the effect of negative income shocks on discount rates; second, comparing the "positive income shock" group to the "always rich" group after the income shock identifies the effect of positive income shocks. Crucially, the two groups being compared have identical wealth levels after the income shock, thus enabling us to compare the effect of income shocks on preferences without confounds from different wealth levels.

Descriptive statistics of the dependent measures by treatment group are shown in Table 4, and results from estimating equation 1 with our different measures of discounting as outcome variables are shown in Table **??**. It can be seen that participants in the "negative income shock" group exhibit greater post-shock discounting than participants in the "always poor" group. Specifically, the mean indifference point between tomorrow and 6 months and one day, the annualized discount rate over this time horizon, and the index for decreasing impatience in this group differ significantly from the "always poor" group, in the direction of greater discounting. Crucially, the wealth levels of these two groups are identical at the time of testing, and thus any differences in discounting must be attributed to the negative income shock per se, rather than to wealth differences. This effect constitutes a selective increase in decreasing impatience: discounting is increased over short time horizons (tomorrow vs. 6 months and 1 day), but not longer time horizons (tomorrow vs. 12 months and 1 day) or time periods in the future (6 months vs. 12 months and 1 day). This is reflected in the selective effect on the mean indifference point and discount rate between tomorrow and 6 months and one day, but not between tomorrow and 12 months and one day, or 6 months and one day and 12 months and one day; accordingly, we find a significant effect of negative income shocks on the index of decreasing impatience.

Is this effect of income shocks on discount rates specific to negative income shocks, or does it also obtain for positive income shocks? The crucial comparison to address this question is between the "always rich" and "positive shock" groups, since the participants in these

groups have identical wealth levels at the time when discount rates are elicited. Table **??** shows an F-test comparing the coefficients on the ALWAYS RICH and POSITIVE SHOCK dummies on our measures of discounting. We find that positive income shocks *decrease* discounting in the distant future, i.e. participants in the "positive shock" condition have a higher indifference point and lower annualized discount rate in tradeoffs between 6 months and one day and 12 months and one day. Thus, discount rates appear to be increased by negative income shocks, and decreased by positive income shocks.

To assess whether persistent low income affected discount rates, we can compare the "always poor" and "always rich" groups. In Table **??**, this amounts to testing whether the coefficient on the ALWAYS RICH dummy is significantly different from zero. We find that indifference points are somewhat lower, and hence discount rates somewhat higher, for the "always rich" group; however, this effect only reaches significance at the 10% level over the time horizon 6 months and one day vs. 12 months and one day. We therefore conclude that persistently low income only weakly affects discount rates.

Effect of Income Shocks and Income Differences on Effort Provision, Social Preferences, and Reservation Wages

We next asked whether income shocks or income differences also affect effort provision, social preferences, or reservation wages. Effort provision was measured by the number of correctly counted tables in Periods 16 and 17, i.e. after the income shock but before the BDM auction. Social preferences were measured with the Ring Task described above. Reservation wages were measured with the BDM auction also described above; participants were given the opportunity to make a bid in a BDM auction on the right to play the real effort task for a further 8 periods after the end of the behavioral tasks that followed period 17. The bid made by participants is an incentive-compatible estimate of the true value to each participant of playing a further 8 periods, and is thus a proxy for the reservation wage of our participants. The results are shown in Table **??**. It can be seen that neither income shocks nor persistent differences in wealth affected effort provision, social preferences, or reservation wages; we only observe a weak negative effect of positive income shocks on effort provision, suggesting that participants who received positive income shocks are less motivated to earn money in subsequent periods because of the sudden windfall gain.

Effect of Income Shocks and Income Differences on Psychological States

Finally, we asked whether the effect of negative income shocks on discount rates might be mediated through effects of the negative income shock on psychological outcomes. We therefore computed the after-before difference of participants' responses on the five visual analog scale questions, and the after-before differences in positive and negative affect as measured by the PANAS scale. The VAS questions elicited self-reported stress, locus of control, optimism, self-esteem, and the degree to which participants thought government should take responsibility for ensuring welfare. The results of OLS regressions of these variables on income shocks are shown in Table **??**. We observe a weak negative effect of negative income shocks on self-reported stress. For positive income shocks, we find and a significant negative effect on self-reported stress levels, along with a weak decrease in negative affect as measured by PANAS, and a significant increase in optimism. We discuss these findings further below.

4 Discussion

The purpose of this study was to test whether income shocks affect preferences. It has already been shown that a correlation exists between income levels and discount rates; in particular, poor people tend to be more impatient than rich people (Lawrance,1991; Sullivan, 2011; Pender & Walker, 1990; Yesuf & Bluffstone, 2008; but see Cagetti, 2003; Gourinchas & Parker, 2002; Stephens & Krupka, 2006; Ogaki & Atkeson, 1997). These studies suffer from the familiar correlation-causality problem: it remains unclear whether poverty actually causes changes in discount rates; in addition, it is not clear to what extent observed differences in discounting behavior actually reflect differences in preferences, or whether they may instead reflect actual or perceived environmental constraints in conditions of poverty. The former question has been addressed to some extent by studies using rainfall data as a source of exogenous variation in income, allowing identification of a causal effect from wealth to discounting behavior (Tanaka et al., 2010; di Falco et al., 2011). However, the second problem persists: it remains unclear whether these observed differences in behavior are reflective of preferences or beliefs. To address this question, we conducted a laboratory experiment in which subjects receive either positive or negative income shocks; crucially, after the shock they have the same level of wealth as a control group that did not receive a shock, allowing a comparison of discount rates across groups which differ only in whether or not they received a shock, not in their levels of wealth. We find that negative income

shocks increase discount rates; this effect is only observed over short time horizons, implying an increase in decreasing impatience (Epper et al., 2011; Prelec et al., 2004). Conversely, positive income shocks weakly decrease discount rates. Together, these findings suggest that discount rates are affected by income shocks; persistently low or high incomes do not affect discounting.

We find no strong effects of income shocks on effort provision, social preferences, and reservation wages in a BDM auction. In terms of psychological outcomes such as positive and negative affect, perceived self-control, optimism, self-esteem, and opinions on government responsibility, we find a weak effect of negative income shocks on self-reported stress, and of positive income shocks on stress, optimim, and negative affect. The weakly significant negative effect of negative income shocks on stress (i.e. participants in the "negative income shock" group showed somewhat lower self-reported levels of stress after than before the shock in comparison to the control group) is somewhat surprising; it is possible that it may reflect a coping strategy to deal with the stressful experience of having received a substantial negative income shock, or that it reflects relief over the resolution of uncertainty that participants experienced when they received the negative income shock.

Less surprising is the highly significant negative effect of positive income shocks on self-reported stress and optimism, and the weaker effect on negative affect. These findings are plausible in light of the fact that subjects are likely to have perceived the positive income shock as a pleasant event.

Our findings also distinguish between two alternative accounts of the effect of reference points on time and risk preferences (Kőszegi & Rabin, 2006, 2007). The first of these is that the negative income shock put participants below their reference point, which was built up during the real effort task. It has been argued that being below the reference point induces risk-seeking behavior (Kahneman & Tversky, 1979); this would predict a decrease in present-focused discount rates because temporally remote outcomes are riskier. Another account by Thaler & Johnson (1990), in contrast, suggests that prior losses introduce risk-averse behavior; this effect in turn should induce present-focused discount rates. Our study therefore contributes to the literature on the effect of reference points on preferences by showing that the second of these accounts is more plausible.

Together, our findings suggest that negative income shocks have a direct effect on economic preferences; in particular, they increase discount rates, and particularly decreasing impatience. It is widely held that humans exhibit more decreasing impatience than is optimal for their own long-run welfare (Laibson, 1997; Prelec, 2004). The mechanism we present here suggests a feedback loop that may account for some of this effect. In particular, if falling into poverty leads to increases in discount rates, then this effect is likely to perpetuate poverty

by leading to imprudent intertemporal decisions.

References

[1] Anderhub, V., Güth, W., Gneezy, U., & Sonsino, D. (2001). On the Interaction of Risk and Time Preferences: An Experimental Study. German Economic Review, 2(3), 239-253. doi:10.1111/1468-0475.00036

[2] Andersen, S., Harrison, G. W., Lau, M. I., & Rutström, E. E. (n.d.). Eliciting Risk and Time Preferences. SSRN eLibrary. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=872325

[3] Banerjee, A. V. (2001). Contracting Constraints, Credit Markets and Economic Development. SSRN eLibrary. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=306990

[4] Bardhan, P., 1996. Research on Poverty and Development Twenty Years after Redistribution with Growth. Proceeding of the Annual World Bank Conference on Development Economics 1995. Supplement to the World Bank Economic Review and the World Bank Research Observer, pp. 59–72.

[5] Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. Behavioral Science, 9(3), 226-232.

[6] Becker, G. S., & Mulligan, C. B. (1997). The Endogenous Determination of Time Preference. The Quarterly Journal of Economics, 112(3), 729 -758. doi:10.1162/003355397555334

[7] Benzion, U., Rapoport, A., & Yagil, J. (1989). Discount Rates Inferred from Decisions: An Experimental Study. Management Science, 35(3), 270 -284. doi:10.1287/mnsc.35.3.270

[8] Bernard, T., Dercon, S., & Taffesse, A.S. (2011). Beyond Fatalism - An Empirical Investigation of Self-Efficacy and Aspirations Failure in Ethiopia. CSAE Working Paper, Oxford.

[9] Bertrand, M., Mullainathan, S., & Shafir, E. (2004). A behavioral-economics view of poverty. American Economic Review: Papers and Proceedings, 94(2), 419-423.

[10] Buckley, N., Chan, K.S., Chowhan, J., Mestelman, S., & Shehata, M. (2001). Value Orientations, Income and Displacement Effects, and Voluntary Contributions. Experimental Economics, 4, 183-195.

[11] Cagetti, M. (2003). Wealth Accumulation Over the Life Cycle and Precautionary Savings. Journal of Business and Economic Statistics, 21(3), 339-353. doi:10.1198/073500103288619007

[12] Chakrabarty, D. (2000). Poverty Traps and Growth in a Model of Endogenous Time Preference. SSRN eLibrary. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=248476

[13] Dasgupta, P. (1997). Environmental and resource economics in the world of the poor. Resources for the Future.

[14] Di Falco, S., Damon,, M., & Kohlin, G. (2011). Environmental Shocks, Rates of Time Preference and Conservation: A Behavioural Dimension of Poverty Traps? Working paper.

[15] Drichoutis, A., & Nayga, R. (2010). Eliciting risk and time preferences under induced mood states. Retrieved from http://mpra.ub.uni-muenchen.de/28338/

[16] Eckel, C., Johnson, C., & Montmarquette, C. (2004). Saving decisions of the working poor: Short- and long-term horizons. In Research in experimental economics, 10, 219–260.

[17] Epper, T., Fehr-Duda, H., & Bruhin, A. (2011). Viewing the future through a warped lens: Why uncertainty generates hyperbolic discounting. Journal of Risk and Uncertainty, 43(3), 169-203. doi:10.1007/s11166-011-9129-x

[18] Frederick, S., Loewenstein, G., & O'donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. Journal of Economic Literature, 40(2), 351-401. doi:10.1257/002205102320161311

[19] Gourinchas, P., & Parker, J. A. (2002). Consumption Over the Life Cycle. Econometrica, 70(1), 47-89. doi:10.1111/1468-0262.00269

[20] Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. Psychological Bulletin, 130(5), 769-792. doi:10.1037/0033-2909.130.5.769

[21] Gunnar Kohlin, Salvatore Di Falco, & Luca Di Corato. (2011). Environmental Shocks and Rates of Time Preference: A Behavioural Dimension of Poverty Traps? Gothenburg University. Retrieved from http://www.webmeets.com/EAERE/2011/prog/viewpaper.asp?pid=914

[22] Hall, Crystal. 2008. \Decisions under Poverty: A Behavioral Perspective on the Decision Making of the Poor." Ph.D. dissertation, Princeton University.

[23] Harris, C., & Laibson, D. (2001). Dynamic Choices of Hyperbolic Consumers. Econometrica, 69(4), 935-957.

[24] Harrison, G. W., Lau, M. I., Rutstrom, E. E., & Williams, M. B. (2002). Estimating Individual Discount Rates in Denmark: A Field Experiment (Artefactual Field Experiments). The Field Experiments Website. Retrieved from http://ideas.repec.org/p/feb/artefa/0053.html

[25] Hausman, J. A. (1979). Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables. The Bell Journal of Economics, 10(1), 33-54. doi:10.2307/3003318

[26] Hershey, J. C., & Schoemaker, P. J. H. (1980). Risk Taking and Problem Context in the Domain of Losses: An Expected Utility Analysis. The Journal of Risk and Insurance, 47(1), 111-132. doi:10.2307/252685

[27] Hershey, J. C., Kunreuther, H. C., & Schoemaker, P. J. H. (1982). Sources of Bias in Assessment Procedures for Utility Functions. Management Science, 28(8), 936-954.

[28] Hill, M. N., McLaughlin, R. J., Pan, B., Fitzgerald, M. L., Roberts, C. J., Lee, T. T.-Y., Karatsoreos, I. N., et al. (2011). Recruitment of prefrontal cortical endocannabinoid signaling by glucocorticoids contributes to termination of the stress response. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 31(29), 10506-10515. doi:10.1523/JNEUROSCI.0496-11.2011

[29] Holden, S.T., Shiferaw, B., & Wik, M. (1998). Poverty, credit constraints, and time preferences: of relevance for environmental policy? Environment and Development Economics, 3, 105-130.

[30] Ifcher, J., & Zarghamee, H. (2008). Happiness and Time Preference: The Effect of Positive Affect in a Random- Assignment Experiment. SSRN eLibrary. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1491247

[31] Joëls, M, & de Kloet, E. R. (1990). Mineralocorticoid receptor-mediated changes in membrane properties of rat CA1 pyramidal neurons in vitro. Proceedings of the National Academy of Sciences of the United States of America, 87(12), 4495-4498.

[32] Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. Nature Neuroscience, 10(12), 1625-1633. doi:10.1038/nn2007

[33] Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. Econometrica, 47(2), 263-291. doi:10.2307/1914185

[34] Karst, H., Berger, S., Turiault, M., Tronche, F., Schutz, G., & Joels, M. (2005). Mineralocorticoid receptors are indispensable for nongenomic modulation of hippocampal

glutamate transmission by corticosterone. Proceedings of the National Academy of Sciences of the United States of America, 102, 19204-19207.

[35] Karst, Henk, Berger, S., Erdmann, G., Schütz, G., & Joëls, M. (2010). Metaplasticity of amygdalar responses to the stress hormone corticosterone. Proceedings of the National Academy of Sciences of the United States of America, 107(32), 14449-14454. doi:10.1073/pnas.0914381107

[36] Kőszegi, B., & Rabin, M. (2006). A Model of Reference-Dependent Preferences. The Quarterly Journal of Economics, 121(4), 1133 -1165. doi:10.1093/qje/121.4.1133

[37] Kőszegi, B., & Rabin, M. (2007). Reference-Dependent Risk Attitudes. The American Economic Review, 97(4), 1047-1073. doi:10.1257/000282807783286829

[38] Krawietz, W., Poppert, D., Erdmann, E., Glossmann, H., Struck, C. J., & Konrad, C. (1976). Beta-adrenergic receptors in guinea-pig myocardial tissue. Naunyn-Schmiedeberg's Archives of Pharmacology, 295(3), 215-224.

[39] Laajaj, Rachid (2012). Closing the Eyes on a Gloomy Future: Psychological Causes and Economic Consequences. Agricultural and Applied Economics Association, 2012 Annual Meeting, Seattle, Washington.

[40] Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. The Quarterly Journal of Economics, 112(2), 443 -478. doi:10.1162/003355397555253

[41] Laibson, D., Repetto, A., & Tobacman, J. (1998). Self-Control and Saving for Retirement. Brookings Papers on Economic Activity, 1, 91–196.

[42] Laibson, D., Repetto, A., & Tobacman, J. (2007). Estimating Discount Functions with Consumption Choices over the Lifecycle (Economics Series Working Papers). University of Oxford, Department of Economics. Retrieved from http://ideas.repec.org/p/oxf/wpaper/341.html

[43] Lawrance, E. C. (1991). Poverty and the Rate of Time Preference: Evidence from Panel Data. Journal of Political Economy, 99(1), 54-77.

[44] Leigh, J. P. (1986). Accounting for Tastes: Correlates of Risk and Time Preferences. Journal of Post Keynesian Economics, 9(1), 17-31.

[45] Loewenstein, G. (1996). Out of Control: Visceral Influences on Behavior. Organizational Behavior and Human Decision Processes, 65(3), 272-292.

[46] Loewenstein, G. (2000). Emotions in Economic Theory and Economic Behavior. The American Economic Review, 90(2), 426-432.

[47] Mazur, J. E. (1988). Estimation of indifference points with an adjusting-delay procedure. Journal of the Experimental Analysis of Behavior, 49(1), 37-47. doi:10.1901/jeab.1988.49-37

[48] Murphree, M. (1993). Communal Land Wildlife Resources and Rural District Council Revenues. CASS, University of Zimbabwe, Harare.

[49] Myerson, J., Green, L., Scott Hanson, J., Holt, D. D., & Estle, S. J. (2003). Discounting delayed and probabilistic rewards: Processes and traits. Journal of Economic Psychology, 24(5), 619-635. doi:10.1016/S0167-4870(03)00005-9

[50] Nguyen, Q., & Leung, P. (2010). How Nurture Can Shape Preferences: An Experimental Study on Risk Preferences of Vietnamese Fishers. Environment and Development Economics, 15(05), 609-631. doi:10.1017/S1355770X10000203

[51] Ogaki, M., & Atkeson, A. (1997). Rate Of Time Preference, Intertemporal Elasticity Of Substitution, And Level Of Wealth, 79(4), 564-572.

[52] Pagiola, S. (1996). Price policy and returns to soil conservation in semi-arid Kenya, 8(3), 225-271.

[53] Patton, J. H., Stanford, M. S., & Barratt, E. S. (1995). Factor structure of the Barratt impulsiveness scale. Journal of Clinical Psychology, 51(6), 768-774.

[54] Payne, J. W., Laughhunn, D. J., & Crum, R. (1980). Translation of Gambles and Aspiration Level Effects in Risky Choice Behavior. Management Science, 26(10), 1039-1060.

[55] Pender, J. L., & Walker, T. S. (1990). Experimental Measurement of Time Preferences in Rural India (No. 97). ICRISAT, Economics Group Progress Report. Andhra Pradesh: ICRISAT: ICRISAT.

[56] Prelec, D. (2004). Decreasing Impatience: A Criterion for Non-stationary Time Preference and "Hyperbolic" Discounting. Scandinavian Journal of Economics, 106, 511–532.

[57] Rachlin, H., Raineri, A., & Cross, D. (1991). Subjective probability and delay. Journal of the Experimental Analysis of Behavior, 55(2), 233-244. doi:10.1901/jeab.1991.55-233

[58] Raeva, D., Mittone, L., & Schwarzbach, J. (2010). Regret now, take it now: On the role of experienced regret in intertemporal choice, 31(4), 634-642.

[59] Richard, T. (1981). Some empirical evidence on dynamic inconsistency. Economics Letters, 8(3), 201-207. doi:10.1016/0165-1765(81)90067-7

[60] Rosenzweig, M. R. (1995). Why Are There Returns to Schooling? The American Economic Review, 85(2), 153-158.

[61] Sarah, L. (1997). The environment and the ethics of discounting: An empirical analysis. Ecological Economics, 20(1), 71-82. doi:10.1016/S0921-8009(96)00073-0

[62] Shamosh, N. A., & Gray, J. R. (2008). Delay discounting and intelligence: A meta-analysis. Intelligence, 36(4), 289-305. doi:10.1016/j.intell.2007.09.004

[63] Shamosh, N. A., Deyoung, C. G., Green, A. E., Reis, D. L., Johnson, M. R., Conway, A. R. A., Engle, R. W., et al. (2008). Individual differences in delay discounting: relation to intelligence, working memory, and anterior prefrontal cortex. Psychological Science, 19(9), 904-911. doi:10.1111/j.1467-9280.2008.02175.x

[64] Shiferaw, B., & Holden, S. T. (2001). Farm-Level Benefits to Investments for Mitigating Land Degradation: Empirical Evidence from Ethiopia. Environment and Development Economics, 6(03), 335-358. doi:10.1017/S1355770X01000195

[65] Slovic, P., Fischhoff, B., & Lichtenstein, S. (1982). Response mode, framing, and information-processing effects in risk assessment. New Directions for Methodology of Social and Behavioral Science: Question Framing and Response Consistency, 21-36.

[66] Stephens, M., & Krupka,, E. L. (2006). Subjective Discount Rates and Household Behavior.

[67] Sullivan, K. A. (2011) The effect of risk, time preference, and poverty on the impacts of forest tenure reform in China. ETD Collection for University of Rhode Island. Paper AAI3464886. http://digitalcommons.uri.edu/dissertations/AAI3464886

[68] Thaler, R. H., & Johnson, E. J. (1990). Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice. Management Science, 36(6), 643-660.

[69] van de Walle, D. (2000). Are returns to investment lower for the poor? Human and physical capital interactions in rural Viet Nam (Policy Research Working Paper Series). The World Bank. Retrieved from http://ideas.repec.org/p/wbk/wbrwps/2425.html

[70] van Stegeren, A., Rohleder, N., Everaerd, W., & Wolf, O. T. (2006). Salivary alpha amylase as marker for adrenergic activity during stress: effect of betablockade. Psychoneuroendocrinology, 31(1), 137-141. doi:10.1016/j.psyneuen.2005.05.012

[71] Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: the PANAS scales. Journal of Personality and Social Psychology, 54(6), 1063-1070.

Table 1: Summary of social preference questions

	Option A		Option B	
Question	Payoff to self	Payoff to other	Payoff to self	Payoff to other
1	600	0	550	230
2	550	230	420	420
3	420	420	230	550
4	230	550	0	600
5	600	0	450	150
6	450	150	300	300
7	300	300	150	450
8	150	450	0	600
9	600	0	370	50
10	370	50	180	180
11	180	180	50	370
12	50	370	0	600

Notes: Summary of social preference questions. Social preferences were measured with a modified version of Buckley et al.'s (2001) Ring Task. Participants answered 12 questions, each of which represented a choice between two options, A and B. Each option contained a payoff for the participant and a payoff for a randomly matched other participant.

Figure 1: Cumulative income during real effort tasks, and income shocks

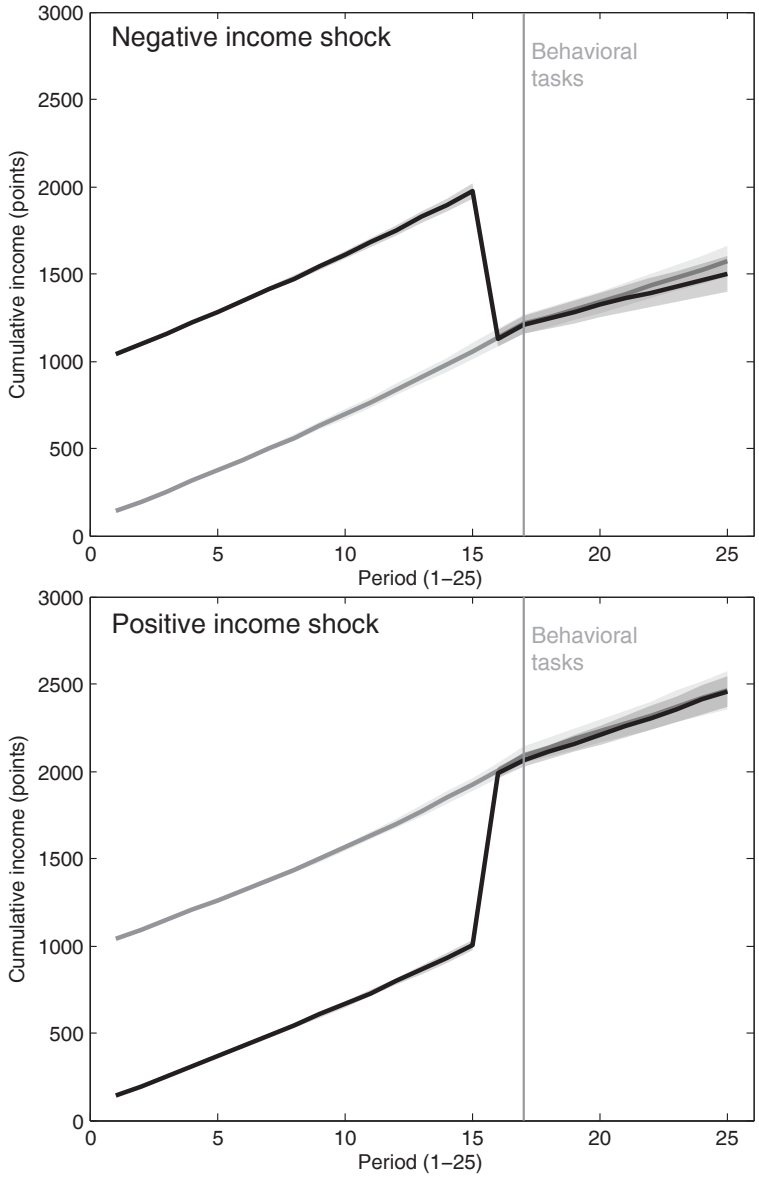

Notes: Cumulative income during real effort tasks and income shocks. The lines show the mean cumulative income across periods for each group. In the top panel, the gray line shows the "always poor" group, the black line the "negative income shock" group and its income shock. In the bottom panel, the gray line shows the "always rich" group, the black line the "positive income shock" group. The shaded areas indicate 1 SEM.

Table 2: Summary statistics

	Negative shock		Positive shock		Always rich		Always poor	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Time preference: indifference points								
Tom. vs. 6 months	17.420	(7.919)	20.790	(7.275)	19.269	(7.853)	21.372	(7.153)
Tom. vs. 12 months	15.443	(8.398)	17.673	(7.601)	17.065	(8.899)	17.496	(8.243)
6 months vs. 12 months	22.259	(7.618)	24.413	(5.022)	21.423	(7.826)	23.678	(5.687)
Time preference: discount rates								
Tom. vs. 6 months	1.366	(1.174)	0.897	(0.892)	1.102	(1.021)	0.825	(0.839)
Tom. vs. 12 months	0.893	(0.832)	0.643	(0.518)	0.833	(0.979)	0.744	(0.828)
6 months vs. 12 months	0.874	(1.483)	0.465	(0.500)	0.879	(1.050)	0.551	(0.619)
Time preference: indices								
Decreasing impatience	0.473	(0.883)	0.254	(0.528)	0.269	(0.737)	0.080	(0.796)
Departure from stationarity	0.492	(1.590)	0.432	(0.844)	0.223	(0.519)	0.274	(0.629)
Effort, social preferences, reservation wage								
Correct in periods 16-17	30.811	(14.508)	28.405	(9.785)	33.784	(16.120)	30.892	(14.132)
Social preferences	1.432	(2.128)	1.432	(1.834)	1.297	(2.222)	1.405	(2.327)
Reservation wage	301.730	(284.977)	311.595	(209.789)	333.676	(298.234)	274.649	(201.927)
Psychological outcomes								
Stress	8.382	(23.211)	0.091	(22.156)	11.257	(20.682)	18.086	(24.827)
Locus of control	−1.735	(17.228)	0.303	(12.890)	0.114	(15.697)	−5.657	(17.083)
Optimism	−7.853	(17.005)	3.030	(17.187)	−1.143	(18.435)	−6.057	(19.969)
Self-esteem	2.382	(12.429)	3.515	(16.782)	−1.829	(14.074)	−1.257	(17.090)
Government responsibility	−0.059	(10.456)	−2.030	(6.361)	0.057	(8.189)	−1.457	(12.955)
PANAS (positive)	−0.315	(0.632)	−0.119	(0.722)	−0.126	(0.686)	−0.066	(0.798)
PANAS (negative)	0.278	(0.407)	0.008	(0.274)	0.118	(0.329)	0.192	(0.482)

Notes: Summary statistics. The columns shown means and standard errors of our variables of interest, separately for the "negative income shock", "positive income shock", "always rich", and "always poor" conditions. The variables of interest are described in detail in the text.

Table 3: Effect of income shocks on discount rates

	Indifference points			Discount rates			Indices	
	(1) Tom. vs. 6 months	(2) Tom. vs. 12 months	(3) 6 months vs. 12 months	(4) Tom. vs. 6 months	(5) Tom. vs. 12 months	(6) 6 months vs. 12 months	(7) Decreasing impatience	(8) Departure from stationarity
Negative shock	−3.953**	−1.991	−1.292	0.546**	0.165	0.326	0.382*	0.221
	(1.795)	(2.001)	(1.561)	(0.247)	(0.205)	(0.284)	(0.199)	(0.310)
Positive shock	−0.764	−0.107	0.881	0.099	−0.088	−0.110	0.187	0.209
	(1.700)	(1.861)	(1.253)	(0.207)	(0.168)	(0.135)	(0.163)	(0.179)
Always rich	−2.211	−0.552	−2.333	0.289	0.097	0.336*	0.192	−0.047
	(1.738)	(2.014)	(1.548)	(0.218)	(0.214)	(0.196)	(0.182)	(0.135)
Constant	16.829***	12.554**	21.774***	1.234**	0.790**	0.539	0.444	0.695
	(4.155)	(4.915)	(4.016)	(0.516)	(0.396)	(0.542)	(0.326)	(0.580)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-tests (p-values)								
Negative vs. positive shock	3.095*	0.980	2.010	3.205*	2.115	2.395	1.153	0.001
	(0.081)	(0.324)	(0.158)	(0.076)	(0.148)	(0.124)	(0.285)	(0.970)
Negative shock vs. Always rich	0.894	0.489	0.368	0.965	0.097	0.001	0.922	0.777
	(0.346)	(0.486)	(0.545)	(0.328)	(0.755)	(0.974)	(0.339)	(0.380)
Positive shock vs. Always rich	0.671	0.052	4.628**	0.701	0.990	5.566**	0.001	2.414
	(0.414)	(0.819)	(0.033)	(0.404)	(0.321)	(0.020)	(0.973)	(0.123)
Observations	148	148	148	148	148	148	148	148
R2	0.071	0.040	0.089	0.061	0.029	0.074	0.042	0.039

Notes: Effect of income shocks on discount rates, OLS regressions. The dependent variables are different measures of discounting; in particular, indifference points (columns (1)-(3)), discount rates (columns (4)-(6)), the measure for decreasing impatience (column (5)), and the measure for departures from stationarity (column (6)). The independent variables are dummies for (1) receiving positive income shocks, (2) negative income shocks, and (3) receiving a large endowment at the beginning of the experiment. The ommited category is the "always poor" condition. Control variables include family income, a dummy for being employed, and a dummy for being currently in debt. White-robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4: Effect of income shocks on effort provision, social preferences, and reservation wages

	(1) Correct in Periods 16-17	(2) Social preferences	(3) Reservation wage
Negative shock	0.087	0.012	47.072
	(3.390)	(0.529)	(60.786)
Positive shock	−2.066	0.079	56.758
	(2.863)	(0.478)	(47.994)
Always rich	2.998	−0.081	59.935
	(3.581)	(0.527)	(56.890)
Constant	30.747***	0.854	128.251
	(8.896)	(1.082)	(125.252)
Controls	Yes	Yes	Yes
F-tests (p-values)			
Negative vs. positive shock	0.511	0.020	0.026
	(0.476)	(0.887)	(0.873)
Negative shock vs. Always rich	0.641	0.034	0.037
	(0.425)	(0.855)	(0.848)
Positive shock vs. Always rich	2.746*	0.111	0.003
	(0.100)	(0.740)	(0.954)
Observations	148	148	148
R2	0.025	0.023	0.065

Notes: Effect of income shocks on effort provision, social preferences, and reservation wages. The dependent variables are: in column (1), the number of tables correctly counted in periods 16 and 17, i.e. after the income shocks but before the other decision-making tasks; in column (2), social preferences on the SVO task; and in column (3), the reservation wage for playing another eight periods after period 17, elicited with a BDM auction. The independent variables are dummies for (1) receiving positive income shocks, (2) negative income shocks, and (3) receiving a large endowment at the beginning of the experiment. The ommited category is the "always poor" condition. Control variables include family income, a dummy for being employed, and a dummy for being currently in debt. White-robust standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 5: Effect of income shocks on psychological outcomes

	(1) Stress	(2) Locus of control	(3) Optimism	(4) Self-esteem	(5) Government responsibility	(6) PANAS (positive)	(7) PANAS (negative)
Negative shock	−10.151*	3.923	−1.758	3.267	2.264	−0.245	0.075
	(5.967)	(4.012)	(4.561)	(3.645)	(2.742)	(0.174)	(0.102)
Positive shock	−19.869***	5.367	9.827**	4.773	0.082	−0.033	−0.214**
	(5.583)	(3.907)	(4.613)	(4.035)	(2.407)	(0.179)	(0.091)
Always rich	−7.173	5.600	4.545	−0.749	1.687	−0.061	−0.081
	(5.365)	(4.009)	(4.459)	(3.757)	(2.587)	(0.169)	(0.094)
Constant	−2.197	−0.303	−2.093	14.036	10.446	0.365	0.009
	(19.666)	(8.950)	(11.505)	(9.145)	(10.080)	(0.461)	(0.226)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-tests (p-values)							
Negative vs.	2.957*	0.161	8.156***	0.176	1.075	1.664	12.855***
positive shock	(0.088)	(0.689)	(0.005)	(0.675)	(0.302)	(0.199)	(0.000)
Negative shock vs.	0.298	0.190	2.324	1.509	0.066	1.341	3.310*
Always rich	(0.586)	(0.664)	(0.130)	(0.222)	(0.797)	(0.249)	(0.071)
Positive shock vs.	6.390**	0.004	1.594	2.142	0.759	0.031	3.633*
Always rich	(0.013)	(0.950)	(0.209)	(0.146)	(0.385)	(0.861)	(0.059)
Observations	137	137	137	137	137	148	148
R2	0.113	0.036	0.129	0.045	0.065	0.041	0.105

Notes: Effect of income shocks on psychological outcomes, OLS regressions. The dependent variables are measures of different psychological states. In columns (1)-(4), subjects responded to visual analog scales (VAS) on a scale of 1-100 to indicate the degree to which they felt stressed (column (1)), in control of their life (column (2)), optimistic (column (3)), and self-esteem (column (4)). In column (5), the dependent variable is agreement with the statement "Government should do more to improve welfare", in opposition to the statement "People should take more responsibility to ensure that they are taken care of". In columns (6) and (7), the dependent variables are positive and negative affect, respectively, as measured by the PANAS scale. The independent variables are dummies for (1) receiving positive income shocks, (2) negative income shocks, and (3) receiving a large endowment at the beginning of the experiment. The ommited category is the "always poor" condition. Control variables include family income, a dummy for being employed, and a dummy for being currently in debt. White-robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.