Negative Income Shocks Increase Discount Rates*

Johannes Haushofer† Ernst Fehr‡

June 10, 2019

Abstract

People with low incomes exhibit higher temporal discounting compared to richer people, i.e., they are more likely to prefer smaller and sooner over larger and later payments. Here we test whether this relationship reflects a causal effect of income on discounting. In doing so, we distinguish whether changes in discounting reflect the experience of a negative shock, or the resulting lower income levels. Participants in a laboratory experiment randomly received different starting endowments, creating "rich" and "poor" groups. All participants then performed a real effort task to earn money, following which subgroups of participants received positive and negative income shocks. Importantly, the magnitude of the shocks and the initial endowments were designed such that they resulted in exactly identical income levels between those groups that did and those that did not receive a shock. This design allows us to identify the effect of income shocks on discounting, while exactly controlling for income levels. We find that negative income shocks lead to an increase in discounting; specifically, they exacerbate present bias, the tendency to overvalue the present relative to the future. Conversely, positive income shocks weakly decrease discounting. In contrast, high vs. low levels of income in the absence of shocks do not affect discounting. The effect of income shocks on discounting cannot be explained by shock-induced stress or negative affect, a desire to break even, or reference point effects. Together, these findings suggest that poverty increases discounting, and that this effect may operate through income shocks rather than income levels.

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

†Department of Psychology, Woodrow Wilson School for Public and International Affairs, and Department of Economics, Princeton University, NJ, USA; National Bureau for Economic Research; and The Busara Center for Behavioral Economics, Nairobi, Kenya. haushofer@princeton.edu

‡Department of Economics, University of Zürich, Switzerland. ernst.fehr@econ.uzh.ch

1. Introduction

Achieving desirable long-run outcomes, such as health and education, often requires making immediate investments that have delayed benefits. People's willingness to make such investments is therefore constrained by temporal discounting, i.e. the subjective de-valuation of the future relative to the present: immediate costs may loom larger than delayed benefits and therefore lead to "impatient" or "present-biased" choices. Poor individuals appear to be particularly prone to such choices: a recent study with 80,000 participants in 76 countries showed a strong positive correlation between income and patience both across and within countries (Dohmen, Enke, Falk, Huffman, & Sunde, 2016).

If this relationship between income and discounting is causal, poverty might perpetuate itself partially through effects on decision-making. Indeed, recent evidence suggests that poverty may affect cognitive processes and decision-making (Mani, Mullainathan, Shafir, & Zhao, 2013; Shah, Mullainathan, & Shafir, 2012), including "over-borrowing" of time from future rounds in a multi-round game show task (Shah et al., 2012; Shah, Mullainathan, & Shafir, 2018). In light of this evidence, and the importance of discounting for economic and social wellbeing, it is perhaps surprising that little is known about whether poverty causally affects discounting.

A possible reason is that answering this question requires exogenous variation in income. Using weather shocks as instrumental variables for income and wealth, Tanaka, Camerer, and Nguyen (2010) and Damon, Di Falco, and Kohlin (2011) show that individuals in Vietnam and Ethiopia who experienced negative income shocks due to droughts have higher discount rates. However, these findings suffer from two problems. First, in real-world settings it is difficult to distinguish between discounting and a subsistence or liquidity constraint: after a drought, people may simply have to put food on the table today, and therefore appear impatient even though the underlying preferences are unchanged (Carvalho, Meier, & Wang, 2016). Second, in these settings it is impossible to disentangle the income effect of the shock from any psychological effects the shock might have. It is thus unclear whether differential discounting in wealth and poverty results from *levels* or *changes*.

Here we use a laboratory study to address these problems with two design features. First, because all payments in our setting are earnings during the experiment, we can rule out that changes in discounting reflect a liquidity constraint. Second, we vary initial income levels and subsequent changes independently, such that we can compare the discount rates of individuals who have exactly equal income *ex post*, but differ in whether or not they recently experienced an income shock. At the beginning of the experiment, participants receive either a high or a low starting endowment, creating "rich" and "poor" groups. Participants then earn additional money in a real effort task, mechanically increasing their income. The motivation behind this element of the design was to give participants a sense of ownership over their experimental income. After about 30 minutes of earning money in the real effort task, we randomly administer downward income shocks to half of the "rich" participants, and upward income shocks to half of the "poor" participants. This creates four groups: participants in the "negative shock" group begin with a high endowment but receive a negative income shock; participants in the "positive shock" group begin with a low endowment but receive a positive income shock; and "always rich" and "always poor" groups of participants begin with high and low endowments, respectively, and do not receive income shocks.

Crucially, the income shocks are structured such that the average income of the shocked group (and its variance) after the shock is precisely identical to the average income (and variance) of one of the groups that never received an income shock. Thus, after the shock, the "negative shock" group has the same average income as the "always poor" group, and the "positive shock" group has the same average income as the "always rich" group. We then administer a standard discounting task, offering participants choices between smaller, sooner versus larger, later payments. This design allows us to compare the discounting behavior of participants whose income *levels* are exactly identical, but who differ in their recent experience of income changes: half of them experienced a positive or negative shock, while the other half did not. We can thus separate the effect of income shocks from the effect of income levels on discounting, in a setting in which subsistence constraints cannot explain behavior.

In studying discounting, we distinguish between long-run impatience on the one hand, and present bias on the other hand. Long-run impatience refers to exponential discounting of future outcomes. Importantly, such discounting is *time consistent*: exponential discounters who prefer one future outcome over another future outcome will do so at any point in time between the present and the sooner of the two outcomes. Present bias, in contrast, refers to a disproportionate value being attached to present outcomes, relative to all future outcomes. Present bias is of central importance in economics and psychology because it predicts *time inconsistency:* present-biased decision-makers will make patient choices between larger future rewards (say in 13 months) and smaller but sooner future rewards (say in 12 months), but then change their choice as the sooner option draws near (Laibson, 1997). Everyday examples include temptation good consumption and procrastination on unpleasant tasks.

We administer a number of additional tasks to rule out confounds and understand the mechanism through which the effect of income shocks on discounting operates. First, to understand if the effect of shocks on discount rates is mediated by a desire to "break even", i.e. recoup the losses incurred through the shock, we measure performance in an additional two rounds of the effort task immediately after the shock, and additionally elicit willingness to pay for the right to complete additional rounds of the effort task. Second, we test whether income shocks affect levels of self-reported stress, optimism, affect, and the stress hormone cortisol.

2. Experimental Design and Analysis

Participants

We recruited 148 male participants form the participant pool of the University of Zürich. Their mean age was 22 years (S.D. 2.47 years). 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 2 hours. Participation was restricted to men because we also measured levels of stress hormones during the experiment, and controlling for ovarian cycle in women is logistically difficult. Participants were native German 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, participants were 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, Clark, & Tellegen, 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. Scales b)–e) were exploratory in nature and are not reported; results are available upon request.

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.

Thus, the "always poor" and "always rich" groups were defined in relative terms, by some half of participants having lower income initially than the other half, and vice-versa. To make it salient to participants that they were either poor or rich relative to others

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

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.

Tasks

Real effort task Participants then participated in a real effort task for 15 periods, which resembled that used by Watson et al. (1988). Each period lasted 2 minutes. The task consisted of counting the number of zeros in a 7 × 5 table of randomly arranged zeros and ones, which was presented on the left side of the screen. The right side of the screen displayed the income variables described above – own income, and maximum, minimum, and average income of all participants. The reason for displaying this information throughout the experiment was to make it continually salient to participants that they were either poorer or richer than others in the study. 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 income of the participant and displayed for 20 seconds in the middle of the screen, again also showing minimum, maximum, and average income. After these 20 seconds, the next period began.

Income shocks Participants played 15 periods of the real effort task, which took about 35 minutes altogether. 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 income levels. Specifically, during the instruction period at the beginning of the study, participants were told that they might experience a change in their income during the real effort task that they would perform. They were told that they would experience either exactly zero or exactly one such income change, but were not told the timing, magnitude, or direction of this change. All participants were told that such sudden changes in income 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 income of this group was equal to the pre-shock average income 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 income of this group was equal to the pre-shock average income of the "always rich" group. Put differently, the two groups "switched places" from the "poor" into the "rich" group, and vice-versa (see Figure 1 for a graphical illustration). This feature of the design allows us to compare the effect of income shocks on economic choice, holding constant levels: 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 income, 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 income. *JH note: You accurately noted that we cannot distinguish the effect of having "lower income than previously" from having "lower income than others in the group". This is true, and we can't solve it in this paper.*

After receiving the income shock, participants were again presented with their updated income and the maximum, minimum, and average income across participants. This information was displayed for one minute to make their new income 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 income 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, participants 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, participants 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 CHF 30, while the sooner smaller reward started at CHF 15 and was then adjusted with a titration method according to the choices the participant made.

Titration is a standard method for identifying discount rates in the discounting literature (Mazur 1988, Green and Myerson 2004, Kable and Glimcher 2007, Rachlin, Raineri, and Cross 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.[1] 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 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. Note that this procedure is unambiguous in identifying an indifference point, in contrast to traditional multiple price lists which can have multiple switch points and therefore multiple candidate indifference points.

This procedure resulted in three individual indifference points for each participant: one comparing tomorrow to 6 months and 1 day; a second comparing tomorrow to 12 months and 1 day; and a third comparing 6 months and 1 day to 12 months and 1 day. To distinguish between long-run impatience and present bias, we proceed as follows. Recall that present bias refers to greater discounting of future outcomes when the sooner outcome is in the present than when both outcomes are in the future. We therefore analyze the impacts of our treatment separately for indifference points between the present and the future (tomorrow vs. 6 months and 1 day, or tomorrow vs. 12 months and 1 day), and indifference points between two future timepoints (6 months and 1 day vs. 12 months and 1 day). This approach allows

[1]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 participant chose CHF 30 in 12 months and 1 day over CHF 15 tomorrow, the next trial would offer the participant a choice between CHF 30 in 12 months and 1 day and CHF 22.50 tomorrow. If the participant 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 participant chose CHF 15 tomorrow over CHF 30 in 12 months and 1 day, the next trial would offer the participant a choice between CHF 7.50 tomorrow and CHF 30 in 12 months and 1 day. If the participant 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.

us to non-parametrically distinguish between present bias and long-run impatience.

For a parametric test of effects on present bias and long-run impatience, we also fit the quasi-hyperbolic discounting model (Laibson, 1997), the most commonly used discounting model in economics. This model has two free parameters: β and δ. The parameter δ is referred to as "long-run impatience" and models the decrease in subjective value over time as an exponential function, with the utility of a payment x after a delay of t months given by $u(x,t) = \delta^t u(x)$. Present bias is modeled by additionally multiplying the subjective value of all outcomes that are *not* in the present with a "present bias" parameter β. The subjective value (utility) of a payment x obtained t periods in the future is therefore given by $u(x,t) = \beta\delta^t u(x)$. In the following we describe how we obtain estimates for β and δ for each individual participant.

Our discounting task identifies indifference points between pairs of timepoints. For instance, if a participant is indifferent between receiving CHF 10 tomorrow and CHF 30 6 months and 1 day from now, we say that a payment of CHF 30 tomorrow would be discounted down to CHF 10 if it was delayed by 6 months. Similarly, if a participant is indifferent between receiving CHF 5 tomorrow and CHF 30 in 12 months and 1 day, we say that a payment of CHF 30 tomorrow would be discounted down to CHF 5 if it was delayed by 12 months. Finally, if a participant is indifferent between receiving CHF 15 6 months and 1 day from now and receiving CHF 30 12 months and 1 day from now, we say that a payment of CHF 30 6 months and 1 day from now would be discounted down to CHF 15 if it was delayed by an additional 6 months, to 12 months and 1 day. We thus have three indifference points per participant, at each of which the utility of receiving the "indifference amount" tomorrow is by definition the same as the utility of receiving CHF 30 at the respective other timepoint (6 months or 12 months). Using linear utility (a reasonable assumption for the stake sizes used here), the quasi-hyperbolic discounting model then allows us to state the following about the discounting parameters δ and β:

$\text{CHF } 10 = \beta\delta^6 \times \text{CHF } 30,$

$\text{CHF } 5 = \beta\delta^{12} \times \text{CHF } 30$

$\beta\delta^6 \times \text{CHF } 15 = \beta\delta^{12} \times \text{CHF } 30$

We now generalize these expressions as follows. We denote the sooner timepoint as t_1 and the later timepoint as t_2. t_1 takes value "0" to refer to "tomorrow", or value "6" to refer to "6 months and 1 day from now. t_2 takes value "6" to refer to "6 months and 1 day from now", and value "12" to refer to "12 months and 1 day from now". We denote the indifference points obtained in the discounting task for the three possible combinations of t_1 and t_2 as x_{t_1,t_2}. In the above examples, $x_{0,6} = \text{CHF } 10$, $x_{0,12} = \text{CHF } 5$, and $x_{6,12} = \text{CHF } 15$.

We next define three normalized indifference points p_{t_1,t_2} between times t_1 and t_2 as

$p_{t_1,t_2} = \frac{x_{t_1,t_2}}{\text{CHF } 30}$. They take value 1 when the participant is fully patient, indicating that the discounted value of the delayed payment is the same as that of the payment at the sooner timepoint. They take value 0 when the participant is fully impatient, indicating that the discounted value of the delayed payment is zero. They take intermediate values when the participant exhibits intermediate levels of discounting.

Substituting the normalized indifference points into the above equations, dividing both sides of each equation by CHF 30, and dividing both sides of the last equation by $\beta\delta^6$, we can then generalize and simplify the three equalities stated above:

$p_{0,6} = \beta\delta^6$

$p_{0,12} = \beta\delta^{12}$

$p_{6,12} = \delta^{12-6} = \delta^6$

We now have a system of 3 equations with 2 unknowns, β and δ. (Recall that the values of $p_{0,6}$, $p_{0,12}$, and $p_{6,12}$ for each participant can be computed by simply dividing the indifference points obtained in the tasks by 30.) Because this system is over-identified (more equations than unkonwns), we cannot compute the values of β and δ algebraically for each participant. We therefore use non-linear least squares estimation to obtain the participant-specific estimates β_i and δ_i that provide the best fit to the data. We run the following non-linear least squares estimation, separately for each participant:

$$p_{t_1,t_2} = \mathbb{1}_{t_1=0}\left[\beta_i\delta_i^{t_2-t_1}\right] + \mathbb{1}_{t_1\neq 0}\left[\delta_i^{t_2-t_1}\right]$$

The symbol $\mathbb{1}$ denotes an indicator function that returns the value 1 when the associated expression (e.g. $t_1 = 0$) is true, and zero otherwise. Thus, this estimation simply models any indifference point between tomorrow ($t_1 = 0$) and some other timepoint t_2 as $\beta_i\delta_i^{t_2-t_1} = \beta_i\delta_i^{t_2}$ (hyperbolic discounting), and the indifference point between two future timepoints (not tomorrow) as $\delta_i^{t_2-t_1}$ (exponential discounting).

The non-linear least squares estimation gives us participant-specific estimates of β_i and δ_i, which we then enter into the regression analysis described below. When $\beta < 1$, the individual is present-biased, with lower values of β indicating increasing present bias. Similarly, $\delta < 1$ indicates long-run impatience, with lower values indicating greater long-run impatience.

Possible serial correlation and order effects in participants' 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 participants 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 the show-up fee of CHF 10 mentioned above, 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. As mentioned above, transaction costs were kept constant by setting the soonest outcome to "tomorrow".

Effort provision and reservation wages We next aimed to test whether the income shocks changed participants' effort provision and reservation wage: participants who just lost a substantial proportion of their income might be more motivated to exert effort in the effort task, and willing to work at a lower wage. We therefore asked participants to play an additional two rounds of the effort task, rounds 16 and 17, and recorded performance in this task. Reservation wages were measured with a BDM auction (Becker, DeGroot, & Marschak, 1964), 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 compared this bid to its own bid. The computer's bid was randomly chosen from a uniform distribution between zero and the expected earnings from a further 8 periods of play, based on performance of each participant in the first 15 periods of the real effort task. If the participants' bid was higher than the computer's bid, the participant paid the computer's bid and completed another 8 rounds of the effort task; if it was lower, the participant paid nothing and did not complete extra rounds of the effort task. The auction was played immediately, and winning participants performed the real effort task for another 8 periods, while the remainder of the participants waited until the winning participants had completed the experiment. The advantage of this type of auction is that it is "incentive-compatible", i.e. it elicits participants' true willingness to pay for playing a further eight periods. To see this, note that if a participant bids below his true valuation, and the computer bids above the participant's bid but below the participant's true valuation, the participant could have played another 8 rounds at the computer's price (which is below his true valuation), but loses the auction and therefore does not have this opportunity. Conversely, if the participant bids more than his true valuation and the computer bids less than the participant but more than the participant's true valuation, the participant wins the auction and now has to pay the computer price to complete the extra 8 periods, but this price lies above his true valuation.

Stress, Affect, and Cortisol levels At the end of the study, participants completed another PANAS questionnaire and the visual-analog scales (see above). Positive affect, negative affect, and stress are analyzed as the difference between the measures obtained after vs. before the study. In addition, we measured levels of the stress hormone cortisol: at the beginning of the session and after the income shock, each participant gave saliva samples which were assayed for cortisol. We analyze the change in cortisol levels from the first to the second sample.

As an exploratory outcome of interest, we also included a social preference task. We found no differences across treatment conditions; detailed results are available upon request. Finally, participants completed a socioeconomic questionnaire and the Barratt Impulsiveness Scale (Patton, Stanford, & Barratt, 1995), and were paid and excused. Heterogeneous treatment effects based on Barratt scores are available upon request.

Statistical analysis

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

$$y_i = \alpha_0 + \alpha_1 \text{NEGATIVE SHOCK}_i + \alpha_2 \text{POSITIVE SHOCK}_i + \alpha_3 \text{ALWAYS RICH}_i + \gamma \mathbf{X_i} + \varepsilon_i \quad (1)$$

where y_i are the outcome variables described above, NEGATIVE SHOCK$_i$, POSITIVE SHOCK$_i$, and ALWAYS RICH$_i$ are dummy variables indicating whether participant 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. The outcome variable y_i represents either measures of discounting, such as indifference points or the individual-level β and δ parameters, or the other measures described above (stress, affect, cortisol, effort provision, etc.). To estimate effects on indifference points involving non-immediate outcomes, we use each participant's indifference point between payments in 6 months and 1 day and 12 months and 1 day as the outcome. To estimate effects on indifference points involving immediate outcomes, we run the regression once for decisions between tomorrow and 6 months and 1 day, and once for decisions between tomorrow and 12 months and 1 day. In principle this approach would yield two sets of coefficient estimates, one from each model. However, our goal is to obtain a single set of coefficients for all decisions involving immediate outcomes. We therefore use seemingly unrelated regression (SUR) to constrain the coefficients to be the same across the two models.

In the tables reporting the results of these regressions, we show the following (combinations of) coefficients: the difference between the "negative shock" and "always poor" group is captured by α_1; the difference between the "positive shock" and "always rich" group is captured by $\alpha_2 - \alpha_3$. Each of these comparisons is equivalent to a t-test with control variables. The regression also allows us to examine if the difference between the negative shock group and the always poor is different from that between the positive shock group and the always rich groups. This difference-in-differences is measured by $\alpha_2 - \alpha_3 - \alpha_1$. This test is identical to the interaction term in a 2x2 ANOVA. Robustness checks in the Supplemental Material report the results without control variables.

3. Results

Evolution of Income

To ascertain that the income shock manipulations worked as intended, we first report the evolution of income levels while performing the real effort task. Figure 1 shows the evolution of income levels as a function of period throughout the experiment. 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 income level in these two groups grew to 1948.38 ± 28.60 (mean ± 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 income 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 income levels through income shocks worked as intended. It can be seen in Figure 1 that the post-shock income levels match exactly those of the "always rich" and "always poor" groups, respectively.

Effect of Income Shocks and Income Differences on Discounting

The main question of this study was whether income shocks affect discounting, while income 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 discounting; 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 income levels after the income shock, thus enabling us to compare the effect of income shocks on preferences without confounds from different income levels.

Figure 2 shows the average indifference points in the discounting task, separately for decisions which involve an immediate option, and decisions in which both options are in the future. Corresponding OLS regressions are shown in Tables S1 (with control variables) and S2 (without control variables). Results are virtually unchanged by the omission of control variables; below we discuss the specifications that include them.

It can be seen that participants in the "negative income shock" group exhibit greater post-shock discounting than participants in the "always poor" group when immediate outcomes are involved: they have lower indifference points ($M = 16.43$, $95\% \, CI = [13.80, 19.07]$) than the always poor group ($M = 19.43$, $95\% \, CI = [17.02, 21.85]$, $t(140) = -2.07$, $p = 0.039$, $d = -0.46$). There is no significant difference between the corresponding indifference points in the "positive income shock" ($M = 19.23$, $95\% \, CI = [16.84, 21.62]$) and "always rich" groups ($M = 18.17$, $95\% \, CI = [15.45, 20.88]$, $t(140) = 0.73$, $p = 0.466$, $d = 0.16$), and the two-way interaction between receiving a shock and the direction of the shock is statistically significant, although barely ($t(140) = 1.99$, $p = 0.047$).

These effects are only seen when decisions involve immediate outcomes: the indifference points for choices between two future outcomes show a similar relationship, but the difference between the "negative income shock" ($M = 22.26$, $95\% \, CI = [19.72, 24.80]$) and "always poor" groups ($M = 23.68$, $95\% \, CI = [21.78, 25.57]$) is not significant ($t(140) = -0.83$, $p = 0.409$, $d = -0.19$). However, when only future outcomes are involved, the "positive income shock" group ($M = 24.41$, $95\% \, CI = [22.74, 26.09]$) has significantly higher indifference points than the "always rich" group ($M = 21.42$, $95\% \, CI = [18.81, 24.03]$, $t(140) = 2.15$, $p = 0.033$, $d = 0.48$). We observe a statistically significant two-way interaction between receiving a shock and the direction of the shock ($t(140) = 2.04$, $p = 0.043$, $d = 0.67$).

Together, these results suggest that discounting is affected by income shocks: we find that negative income shocks increase present bias. At the same time, positive income shocks decrease long-run discounting. To test these effects parametrically, Figure 3 shows the β and δ parameters from the quasi-hyperbolic discounting model (Laibson, 1997), corresponding to present bias and long-run impatience, respectively, and again Tables S1 (with control variables) and S2 (without control variables) show results from the corresponding OLS regressions. We find a strong and statistically significant effect of negative income shocks on present bias: We observe substantially lower β parameters (i.e. more present bias) in the "negative shock" group ($M = 0.75$, $95\% \, CI = [0.65, 0.85]$) compared to the "always poor" group ($M = 0.89$, $95\% \, CI = [0.79, 0.99]$, $t(140) = -1.99$, $p = 0.048$, $d = -0.49$). No such difference is observed for positive income shocks ("positive shock" group $M = 0.86$, $95\% \, CI = [0.76, 0.96]$); "always rich" group $M = 0.90$, $95\% \, CI = [0.81, 0.99]$, $t(140) = -0.76$, $p = 0.45$, $d = -0.17$), although the interaction is not statistically significant ($t(140) = 0.94$, $p = 0.349$). We observe no significant effects of negative shocks on long-run impatience (δ; "negative shock" group $M = 0.92$, $95\% \, CI = [0.86, 0.98]$), "always poor" group $M = 0.93$, $95\% \, CI = [0.88, 0.99]$, $t(140) = -0.43$, $p = 0.732$, $d = -0.09$). These results mirror the fact that, as shown above, negative shocks affect the indifference points for decisions involving immediate outcomes, but not non-immediate outcomes. In contrast, positive income shocks lead to less long-run discounting in the δ parameter ("positive shock" group $M = 0.97$, $95\% \, CI = [0.95, 0.98]$, "always rich" group $M = 0.92$, $95\% \, CI = [0.86, 0.97]$), again mirroring the effect described above for indifference points. However, this effect fails to reach statistical significance at conventional levels ($t(140) = 1.95$, $p = 0.053$, $d = 0.39$). We observe little evidence of an interaction ($t(140) = 1.43$, $p = 0.155$, $d = 0.47$).

To assess whether persistent low income affected discount rates, we can compare the "always poor" and "always rich" groups. We find no significant effects of persistently low income on discount rates (results not shown).

Effect of Income Shocks and Income Differences on Effort Provision and Reservation Wages

We now turn to investigating the mechanisms behind the effect of income shocks on discounting. We begin by asking whether income shocks move participants below the reference point, such that they are motivated to "make up" for the "lost" income on the day of the experiment by preferentially choosing outcomes that are available sooner. We address this question by asking if income shocks also affect effort provision 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. Reservation wages were measured with the BDM auction described above, which yields a measure of the reservation wage for playing additional periods. The results are shown in Figure 4, and corresponding OLS regressions are shown in Tables S3 (with control variables) and S4 (without control variables). It can be seen that neither income shocks nor persistent differences in income affected effort provision or reservation wages; we only observe a weak negative effect of positive income shocks on effort provision ($p = 0.10$), suggesting that participants who received positive income shocks are less motivated to earn money in subsequent periods because of the sudden windfall gain. However, none of the comparisons reach conventional levels of significance ($p > 0.25$).

Effect of Income Shocks and Income Differences on Psychological States and Hormone Levels

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 self-reported stress, positive and negative affect as measured by the PANAS scale, and cortisol levels. The results of these analyses are shown in Figure 5, and corresponding OLS regressions are shown in Tables S5 (with control variables) and S6 (without control variables). We observe a weak negative effect of negative income shocks on self-reported stress, but this effect is not statistically significant at conventional levels. For positive income shocks, we find and a significant negative effect on self-reported stress levels in the specification with controls variables ($p = 0.042$), but not the specification without control variables; we therefore do not interpret this result confidently.

4. Discussion

The purpose of this study was to test whether income shocks affect temporal discounting. Previous work has shown that poor individuals and countries exhibit higher discounting than others (Falk et al. 2018, Lawrance 1991, Sullivan 2011, Pender and Walker 1990, Yesuf and Bluffstone 2008, Stephens and Krupka 2006). However, these studies suffer from the familiar correlation-causality problem: it remains unclear whether poverty actually causes changes in discount rates. Studies that address this problem using instrumental variables have suggested that negative income shocks may increase discount rates (Damon et al., 2011; Tanaka et al., 2010). However, in these studies, it is difficult to distinguish the effect of low income levels from that of negative shocks. 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 constraints, such as subsistence or liquidity constraints (Carvalho et al., 2016).

The core element of our design is that it allows us to disentangle the effect of shocks from an effect of levels, and does not suffer from confounds due to liquidity constraints: by randomly assigning participants to difference income levels and then randomly and independently assigning them to positive and negative shocks, we can compare groups of participants who are otherwise identical but either have different income, or have the same income but have recently experienced vs. not experienced a positive or negative shock.

We find that participants who experienced negative income shocks ("negative shock" group) discount more than those who did not experience such shocks but who have the same income levels ("always poor" group). In contrast, positive income shocks decrease discounting somewhat, although this effect is weaker. Thus, our evidence suggests that negative income shocks increase discount rates, holding constant income levels.

We distinguish two facets of discounting: present bias and long-run impatience. The effect of negative income shocks is specific to present bias, and is not found for long-run impatience. Because present bias predicts time inconsistent behavior, this finding has potential policy relevance: when individuals suffer downward income shocks, they may find it difficult to make time-consistent decisions, e.g. by following through on existing plans. In contrast, positive income shocks decrease long-run impatience, although this result is only statistically significant at conventional levels in the analysis using indifference points and not in that using the quasi-hyperbolic discounting model. This is likely a power issue.

We administer a number of additional tasks to rule out confounds and understand the mechanism through which the effect of income shocks on discounting operates. First, we let participants play an additional two rounds of the effort task immediately after the shock, and find no effect of shocks on participants' effort provision in these rounds. This finding suggests that the effect of shocks on discount rates is not mediated by a desire to "break even", i.e. recoup the losses incurred through the shock. A further piece of evidence in support of this finding is that when we subsequently offer participants an opportunity to buy the right to complete additional rounds of the effort task, their willingness to pay for this right does not differ as a result of the shock.

Second, note also that changes in risk seeking as a result of being in a "loss frame" after the shock cannot account for the effect of shocks on discounting because they would predict the effect to go in the opposite direction: in a loss frame, individuals are more risk seeking, and therefore they should also be more tolerant of the risk associated with choosing delayed outcomes in the discounting task.

Finally, we find a negative effect of both positive and negative income shocks on self-reported stress. Because these effects on stress go in the same direction, while positive and negative shocks affect discounting in opposite directions, it is unlikely that stress is the mechanism that drives the effects of shocks on discounting. Rather, the fact that the effects are similar for both positive and negative shocks suggests that they reflect relief over the resolution of uncertainty that participants experienced when they received a shock.

In sum, our results are unlikely to be due to changes in affect, a desire to break even, or reference point effects. One remaining potential explanation for our findings is that shocks increase participants' perception of background risk, leading them to prefer safer, sooner outcomes. Future research might test this hypothesis.

This study contributes to an emerging literature on the effects of poverty on economic choice. A number of authors have suggested that poverty may affect psychological processes (Haushofer & Fehr, 2014), cognition, and decision-making (Mani et al., 2013; Shah et al., 2012). Indeed, Shah et al. (2012) show that people in a game-show like task "over-borrow" time from future periods to complete current tasks, a behavior which can be thought of as discounting future periods too steeply. More specifically for discounting, two previous studies have shown causal evidence in field settings for an effect of income shocks on discount rates (Tanaka et al. 2010 and Damon et al. 2011). Our findings suggest that these effects may be due to the psychological consequences of the shocks per se, rather than the different income levels resulting from them. A related laboratory study is that by Raeva, Mittone, and Schwarzbach (2010), who experimentally induced "regret" or "rejoicing" in participants by offering them a choice between two lotteries and then revealing both the obtained and forgone payoff. Participants experiencing regret had higher discount rates than controls, while the opposite was true for those experiencing rejoicing. It is possible that a similar regret mechanism was at play in our study, although we explicitly informed participants that they had no control over the income shocks. A further related study is that by Spears (2011), who showed in a lab-in-the-field experiment in India that making choices under conditions of scarcity (having a lower experimental endowment) impaired cognitive control in a handgrip and Stroop task. These results suggest that even levels of "wealth" may have effects on subsequent behavior, in contrast to our findings which indicate an effect of shocks. The outcome measures used by Spears (2011) differ from ours, but cognitive control has frequently been related to hyperbolic discounting in the literature (Shamosh et al. 2008, Shamosh and Gray 2008), and thus it is possible that this effect would extend to discounting.

More broadly, our results complement those of several studies on the effect of induced emotions on discounting. Loewenstein (1996, 2000) first pointed out that in the presence of visceral factors such as rage, people sometimes exhibit extreme discounting of future events. In line with the hypothesis that affect may change behavior, Lerner, Li, and Weber (2013) found that participants exhibited higher discount rates after they had watched sad compared to neutral video clips. These findings are consistent with the view that discount rates respond to experimentally induced affective states, and raise the possibility that the findings we report here were in fact due to such affective changes, but that we did not have sufficient power to detect them with our affect measures.

Together, our findings suggest that negative income shocks have a direct effect on economic preferences; in particular, they increase discounting, particularly present bias. It is widely held that humans exhibit more discounting 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: if falling into poverty leads to increases in discount rates, then this effect may perpetuate poverty by leading to imprudent inter-temporal decisions.

References

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

Carvalho, L. S., Meier, S., & Wang, S. W. (2016). Poverty and economic decision-making: Evidence from changes in financial resources at payday. *American Economic Review*, *106*(2), 260-84.

Damon, M., Di Falco, S., & Kohlin, G. (2011). Environmental Shocks and Rates of Time Preference: A behavioural Dimension of Poverty Traps? *Working Paper*.

Dohmen, T., Enke, B., Falk, A., Huffman, D., & Sunde, U. (2016). *Patience and the Wealth of Nations* (Working Paper). Human Capital and Economic Opportunity Working Group.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., & Sunde, U. (2018). Global Evidence on Economic Preferences. *The Quarterly Journal of Economics*, qjy013.

Green, L., & Myerson, J. (2004). A Discounting Framework for Choice With Delayed and Probabilistic Rewards. *Psychological Bulletin*, *130*(5), 769-792.

Haushofer, J., & Fehr, E. (2014). On the psychology of poverty. *Science*, *344*(6186), 862-867.

Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, *10*(12), 1625-1633.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*, *112*(2), 443-477.

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

Lerner, J. S., Li, Y., & Weber, E. U. (2013). The financial costs of sadness. *Psychological Science*, *24*(1), 72-79.

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

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

Mani, A., Mullainathan, S., Shafir, E., & Zhao, J. (2013). Poverty Impedes Cognitive Function. *Science*, *341*(6149), 976-980.

Mazur, J. E. (1988). Estimation of indifference points with an adjusting-delay procedure. *Journal of the Experimental Analysis of Behavior*, *49*(1), 37-47.

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.

Pender, J. L., & Walker, T. S. (1990). Experimental measurement of time preference in rural India. *Working Paper*.

Prelec, D. (2004). Decreasing Impatience: A Criterion for Non-stationary Time Preference and "Hyperbolic" Discounting. *Scandinavian Journal of Economics*, *106*(3), 511-532.

Rachlin, H., Raineri, A., & Cross, D. (1991). Subjective probability and delay. *Journal of the Experimental Analysis of Behavior*, *55*(2), 233-244.

Raeva, D., Mittone, L., & Schwarzbach, J. (2010). Regret now, take it now: On the role of experienced regret on intertemporal choice. *Journal of Economic Psychology*, *31*, 634-642.

Shah, A. K., Mullainathan, S., & Shafir, E. (2012). Some Consequences of Having Too Little. *Science*, *338*(6107), 682-685.

Shah, A. K., Mullainathan, S., & Shafir, E. (2018). An opportunity for self-replication. *Nature Human Behaviour*, 1.

Shamosh, N. A., Deyoung, C. G., Green, A. E., Reis, D. L., Johnson, M. R., Conway, A. R. A., ... Gray, J. R. (2008). Individual differences in delay discounting: Relation to intelligence, working memory, and anterior prefrontal cortex. *Psychological Science*, *19*(9), 904-911.

Shamosh, N. A., & Gray, J. R. (2008). Delay discounting and intelligence: A meta-analysis. *Intelligence*, *36*(4), 289-305.

Spears, D. (2011). Economic Decision Making in Poverty Depletes Behavioral Control. *B.E. Journal of Economic Analysis and Policy*, *11*(1), 1-44.

Stephens, M., & Krupka, E. (2006). Subjective Discount Rates and Household Behavior. *Working Paper*.

Sullivan, K. (2011). The Effect of Risk, Time Preference, and Poverty on the Impacts of Forest Tenure Reform in China. *Dissertation, University of Rhode Island*.

Tanaka, T., Camerer, C. F., & Nguyen, Q. (2010). Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam. *American Economic Review*, *100*(1), 557-571.

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.

Yesuf, M., & Bluffstone, R. (2008). Wealth and Time Preference in Rural Ethiopia. *Working Paper*.

Figures

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

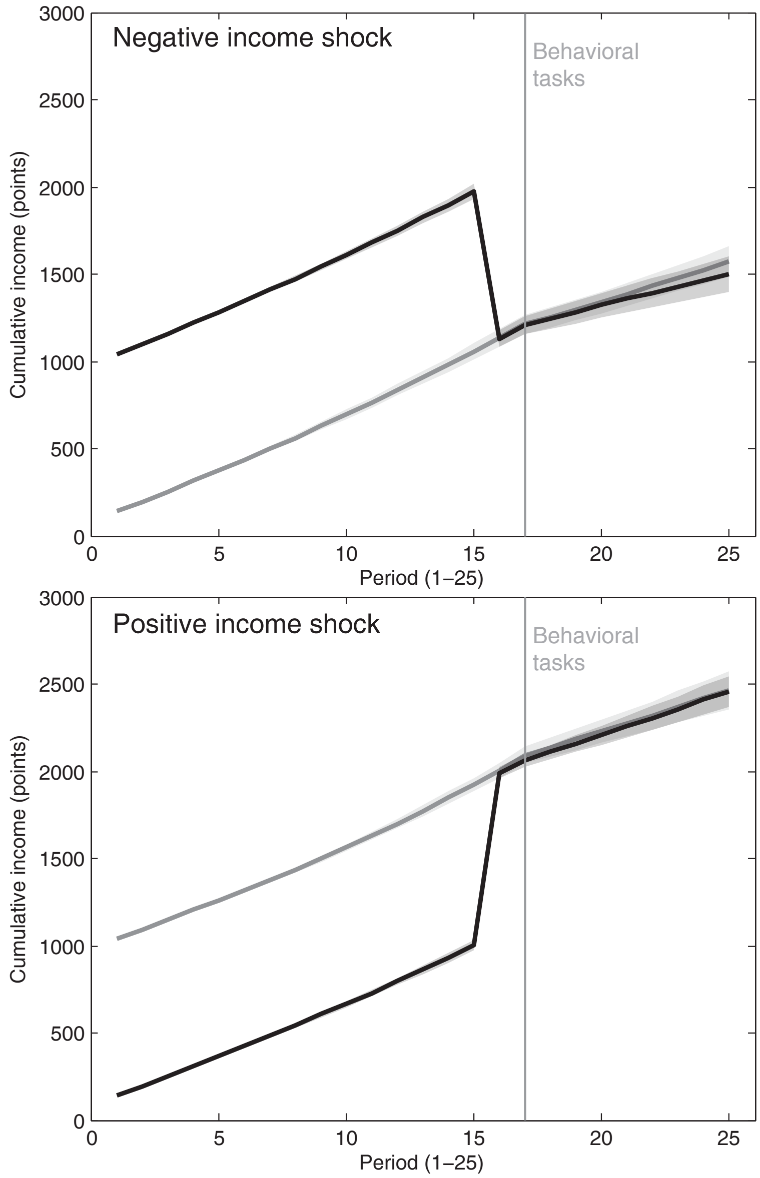

Notes: Cumulative income during real effort tasks and income shocks. The lines show the mean cumulative income across periods for each group. In the top panel, the gray line shows the "always poor" group, the black line the "negative income shock" group and its income shock. In the bottom panel, the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

Figure 2: Effect of Income Shocks on Indifference Points

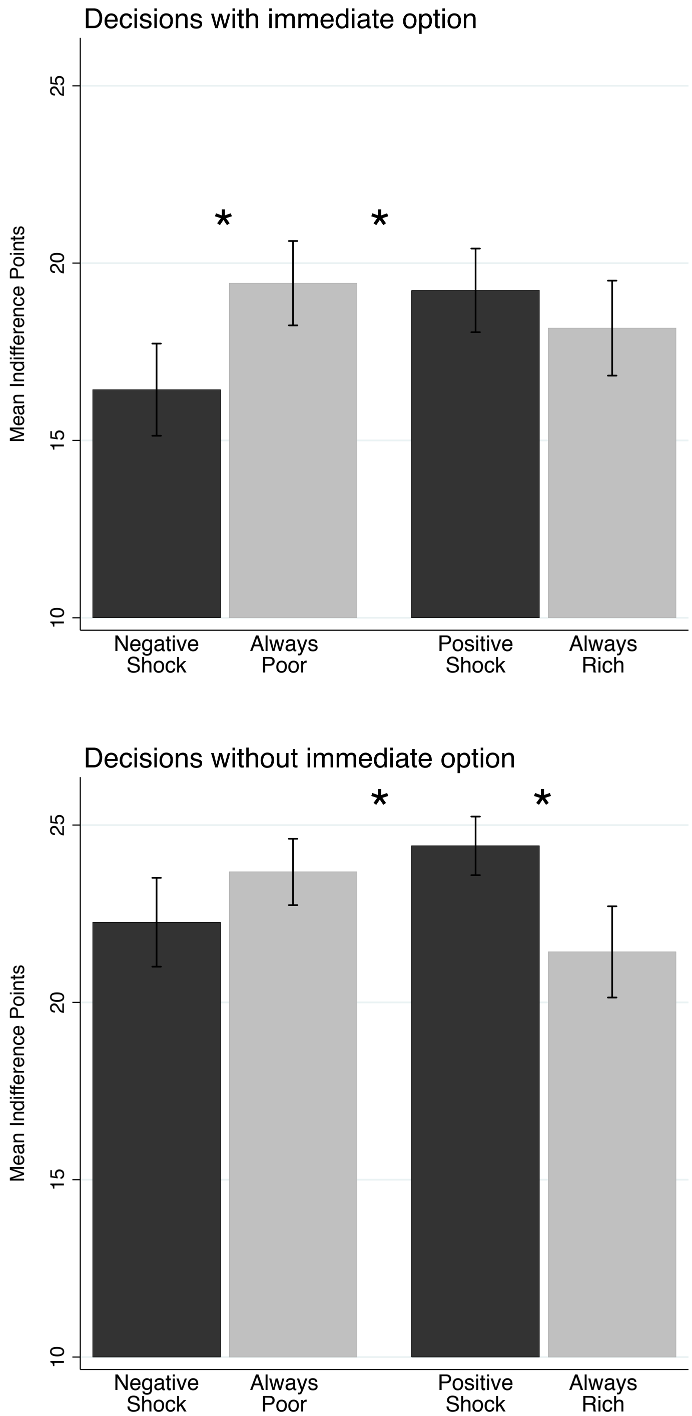

Notes: Mean indifference points across "Negative Shock", "Always Poor", "Positive Shock" and "Always Rich" conditions. The asterisks denote significant differences between conditions based on OLS regressions. Asterisks between pairs of bars reflect significant interaction terms. Error bars indicate 1 SEM. *** $p<0.001$, ** $p<0.01$, * $p<0.05$.

Figure 3: Effect of Income Shocks on Impatience and Present Bias

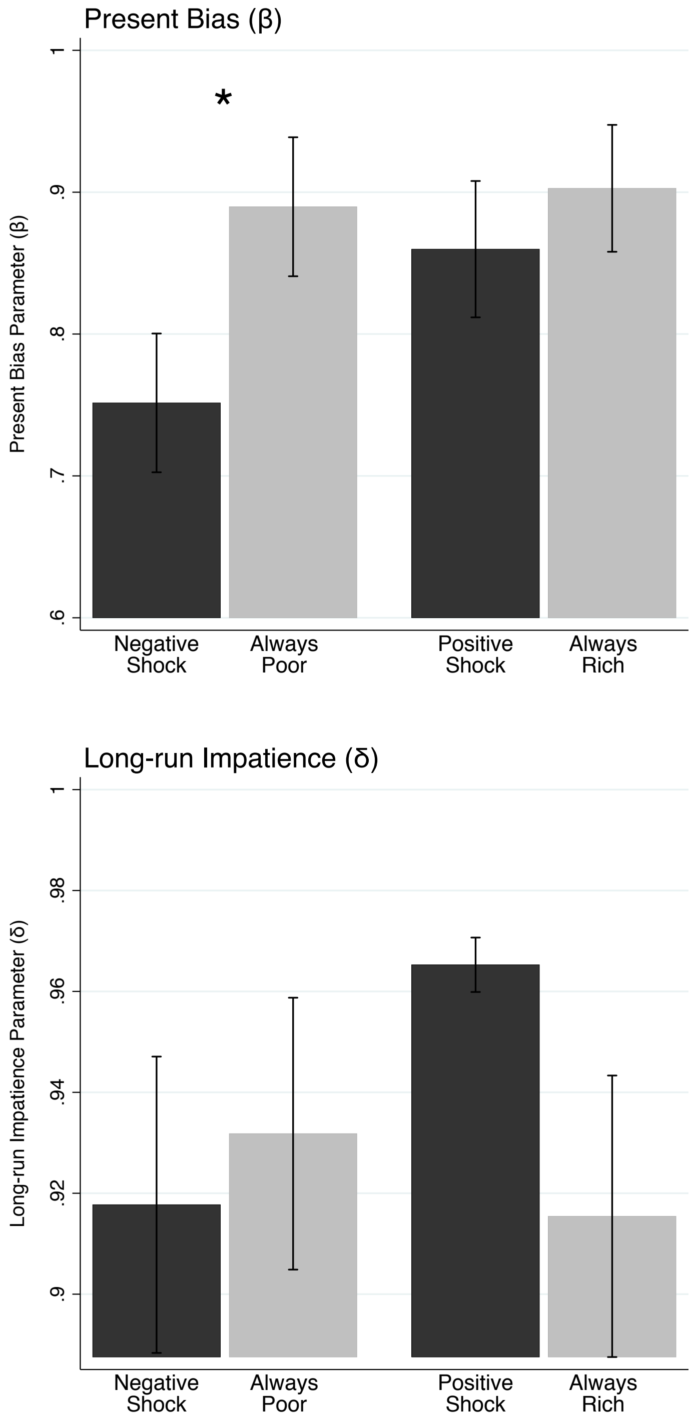

Notes: Mean values of the present bias parameter β and the long-run impatience parameter δ across "Negative Shock", "Always Poor", "Positive Shock" and "Always Rich" conditions. The asterisks denote significant differences between conditions based on OLS regressions. Asterisks between pairs of bars reflect significant interaction terms. Error bars indicate 1 SEM. *** p<0.001, ** p<0.01, * p<0.05.

Figure 4: Effect of income shocks on Effort Provision and Reservation Wages

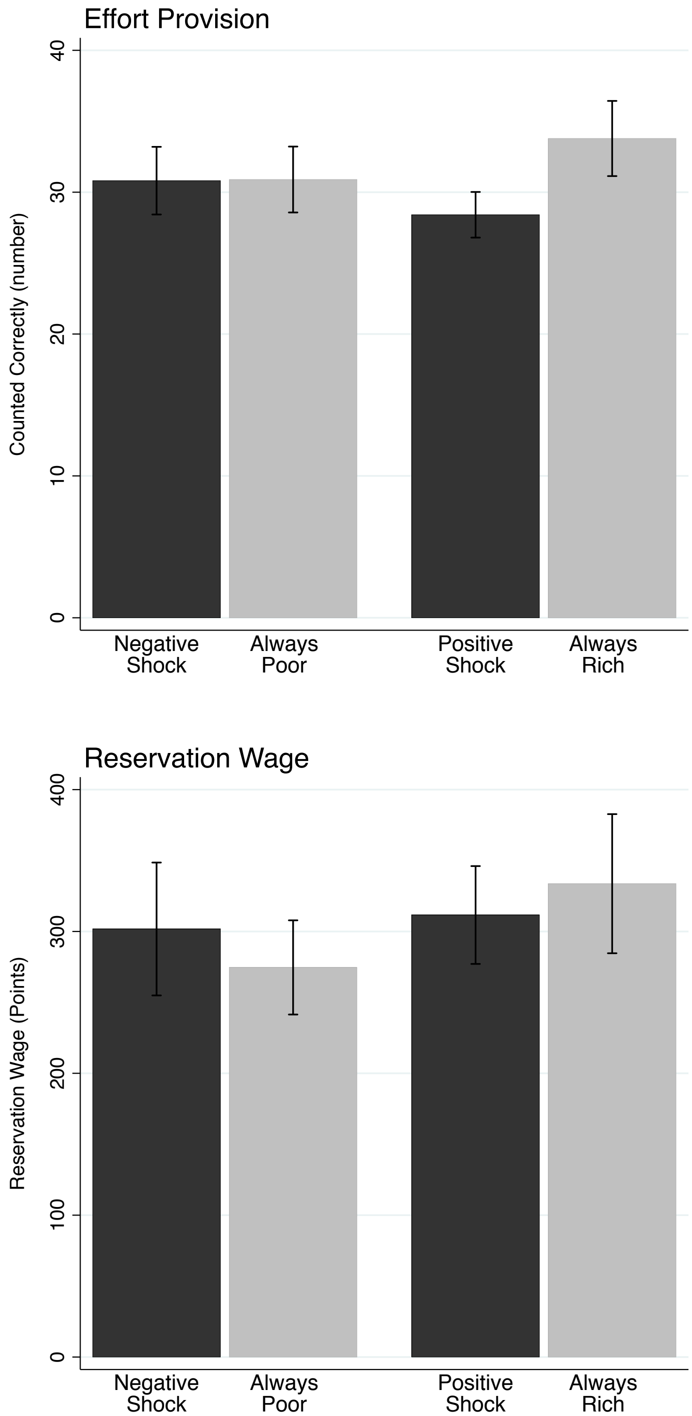

Notes: Mean values of effort provision (measured as numbers of tables completed in the real effort task in periods 16 and 17) and reservation wages (measured as willingness to pay for the opportunity to complete additional periods of the effort task) across “Negative Shock”, “Always Poor”, “Positive Shock” and “Always Rich” conditions. The asterisks denote significant differences between conditions based on OLS regressions. Asterisks between pairs of bars reflect significant interaction terms. Error bars indicate 1 SEM. *** $p<0.001$, ** $p<0.01$, * $p<0.05$.

Figure 5: Effect of income shocks on Stress, Cortisol and Affect

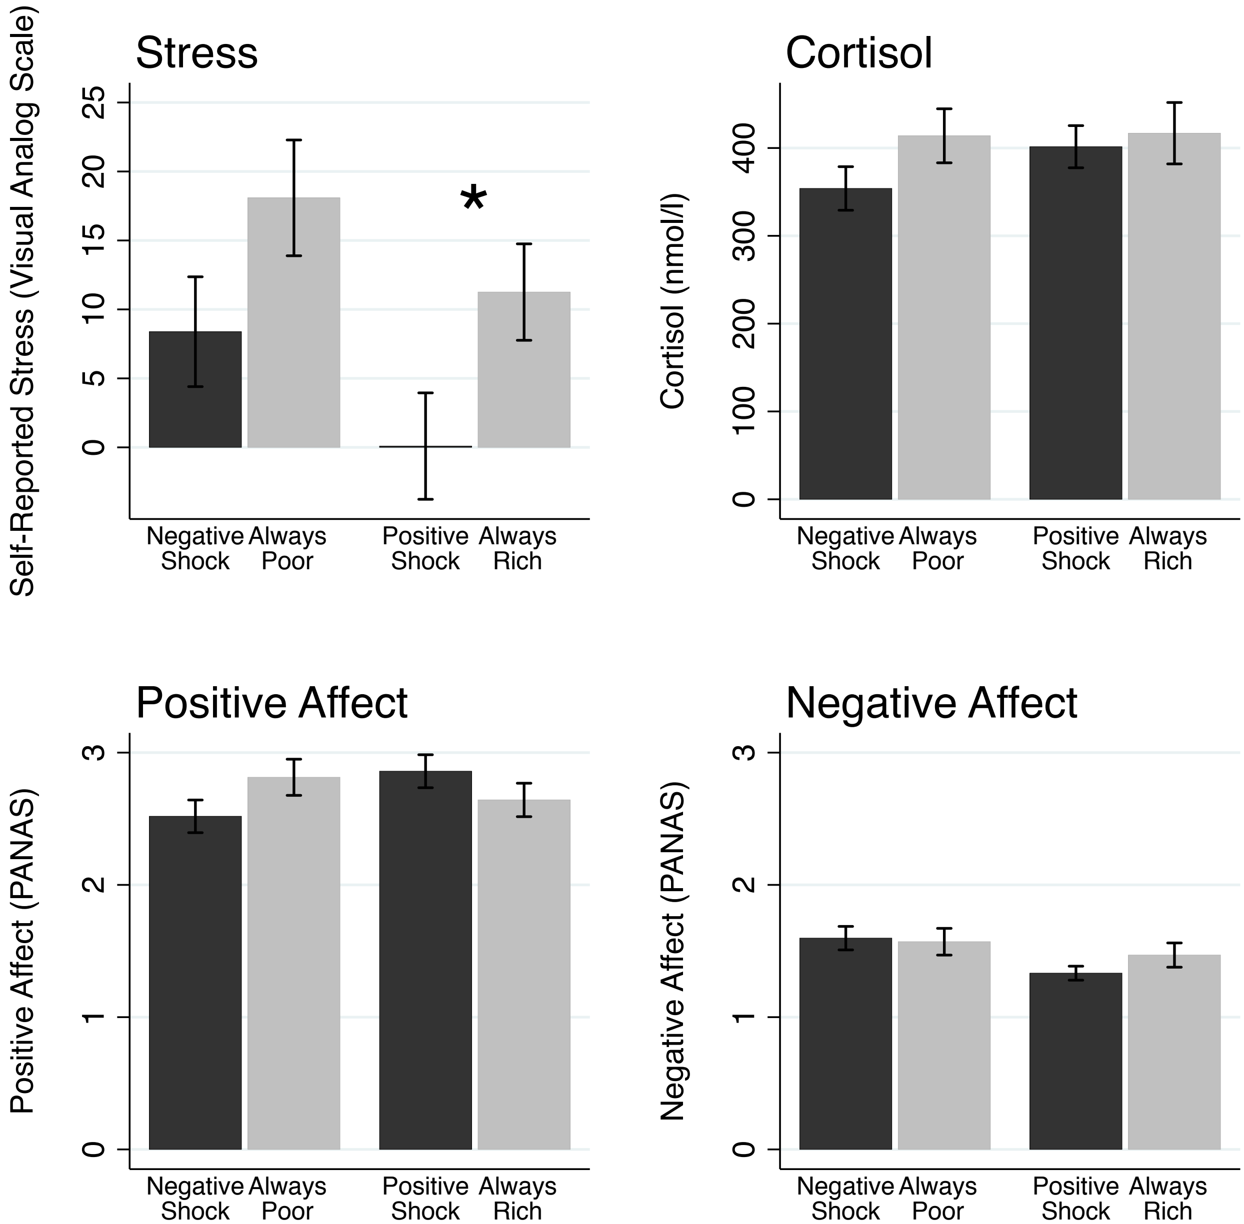

Notes: Mean levels of self-reported stress (measured with a visual analog scale from 0–100), cortisol (nmol/l), and positive and negative affect (measured with the PANAS scale) across "Negative Shock", "Always Poor", "Positive Shock" and "Always Rich" conditions. The asterisks denote significant differences between conditions based on OLS regressions. Asterisks between pairs of bars reflect significant interaction terms. Error bars indicate 1 SEM. *** $p<0.001$, ** $p<0.01$, * $p<0.05$.

Negative Income Shocks Increase Discount Rates: Supplemental Material

Johannes Haushofer[1] Ernst Fehr[2]

[1]Department of Psychology, Woodrow Wilson School for Public and International Affairs, and Department of Economics, Princeton University, NJ, USA; National Bureau for Economic Research; and The Busara Center for Behavioral Economics, Nairobi, Kenya. haushofer@princeton.edu

[2]Department of Economics, University of Zürich, Switzerland. ernst.fehr@econ.uzh.ch

Table S1: Effect of Income Shocks on Discounting

	Indifference points		Quasi-hyperbolic discounting model	
	Decisions with immediate option	Decisions without immediate option	Present Bias (β)	Long-run Impatience (δ)
Negative Shock vs. Always Poor	-3.55^{*}	-1.29	-0.14^{*}	-0.01
	(1.72)	(1.56)	(0.07)	(0.04)
Positive Schock vs. Always Rich	1.25	3.21^{*}	-0.05	0.06
	(1.71)	(1.49)	(0.07)	(0.03)
Interaction	4.78^{*}	4.51^{*}	0.09	0.07
	(2.41)	(2.21)	(0.10)	(0.05)
Observations	148	148	148	148

Notes: Effect of income shocks on indifference points and parameters for present bias and long-run impatience, OLS regressions with control variables. The dependent variables are different measures of discounting; in particular, indifference points (columns (1)-(3)), and the β parameter for present bias and the δ parameter for long-run impatience in the quasi-hyperbolic discounting model (columns (4)-(5)). The estimates in each row represent an effect of interest (coefficient or combination of coefficients), and its robust standard error in parentheses. The estimates are obtained through parameter combinations of our main specification that identify the effect in question, as described in Section 2. The first row shows the difference between the "negative shock" and "always poor" groups; the second row between the "positive shocks" and "always rich" groups; and the third row shows the difference between these two effects. Control variables are included in all regressions and include family income, a dummy for being employed, and a dummy for being currently in debt. *** $p<0.001$, ** $p<0.01$, * $p<0.05$.

Table S2: Effect of Income Shocks on Discounting (No Control Variables)

	Indifference points		Quasi-hyperbolic discounting model	
	Decisions with immediate option	Decisions without immediate option	Present Bias (β)	Long-run Impatience (δ)
Negative Shock vs. Always Poor	−3.56*	−1.42	−0.14*	−0.01
	(1.72)	(1.56)	(0.07)	(0.04)
Positive Schock vs. Always Rich	1.34	2.99	−0.04	0.05
	(1.72)	(1.53)	(0.07)	(0.03)
Interaction	4.90*	4.41*	0.10	0.06
	(2.44)	(2.19)	(0.10)	(0.05)
Observations	148	148	148	148

Notes: Effect of income shocks on indifference points and parameters for present bias and long-run impatience, OLS regressions without control variables. The dependent variables are different measures of discounting; in particular, indifference points (columns (1)-(3)), and the β parameter for present bias and the δ parameter for long-run impatience in the quasi-hyperbolic discounting model (columns (4)-(5)). The estimates in each row represent an effect of interest (coefficient or combination of coefficients), and its robust standard error in parentheses. The estimates are obtained through parameter combinations of our main specification that identify the effect in question, as described in Section 2. The first row shows the difference between the "negative shock" and "always poor" groups; the second row between the "positive shocks" and "always rich" groups; and the third row shows the difference between these two effects. Control variables are not included. *** p<0.001, ** p<0.01, * p<0.05.

Table S3: Effect of Income Shocks on Effort Provision and Reservation Wages

	Effort Provision	Reservation Wage
Negative Shock vs. Always Poor	0.09	47.07
	(3.39)	(60.79)
Positive Schock vs. Always Rich	−5.06	−3.18
	(3.06)	(54.42)
Interaction	−5.15	−50.25
	(4.62)	(81.91)
Observations	148	148

Notes: Effect of income shocks on effort provision and reservation wages, OLS regressions with control variables. The dependent variables are effort provision (column (1)), measured as numbers of tables completed in the real effort task in periods 16 and 17, and reservation wages (column (2)), measured as willingness to pay for the opportunity to complete additional periods of the effort task. The estimates in each row represent an effect of interest (coefficient or combination of coefficients), and its robust standard error in parentheses. The estimates are obtained through parameter combinations of our main specification that identify the effect in question, as described in Section 2. The first row shows the difference between the "negative shock" and "always poor" groups; the second row between the "positive shocks" and "always rich" groups; and the third row shows the difference between these two effects. Control variables are included in all regressions and include family income, a dummy for being employed, and a dummy for being currently in debt. *** $p<0.001$, ** $p<0.01$, * $p<0.05$.

Table S4: Effect of Income Shocks on Effort Provision and Reservation Wages (No Control Variables)

	Effort Provision	Reservation Wage
Negative Shock vs. Always Poor	−0.08	27.08
	(3.33)	(57.42)
Positive Schock vs. Always Rich	−5.38	−22.08
	(3.10)	(59.94)
Interaction	−5.30	−49.16
	(4.55)	(83.01)
Observations	148	148

Notes: Effect of income shocks on effort provision and reservation wages, OLS regressions without control variables. The dependent variables are effort provision (column (1)), measured as numbers of tables completed in the real effort task in periods 16 and 17, and reservation wages (column (2)), measured as willingness to pay for the opportunity to complete additional periods of the effort task. The estimates in each row represent an effect of interest (coefficient or combination of coefficients), and its robust standard error in parentheses. The estimates are obtained through parameter combinations of our main specification that identify the effect in question, as described in Section 2. The first row shows the difference between the "negative shock" and "always poor" groups; the second row between the "positive shocks" and "always rich" groups; and the third row shows the difference between these two effects. Control variables are not included. *** p<0.001, ** p<0.01, * p<0.05.

Table S5: Effect of income shocks on Stress, Cortisol Levels and Positive and Negative Affect

	Stress	Cortisol	Positive Affect	Negative Affect
Negative Shock vs.	−10.15	−64.28	−0.27	0.07
Always Poor	(5.97)	(40.75)	(0.17)	(0.10)
Positive Schock vs.	−12.70*	−10.62	0.10	−0.13
Always Rich	(5.02)	(42.81)	(0.16)	(0.07)
Interaction	−2.54	53.66	0.37	−0.20
	(7.79)	(56.49)	(0.23)	(0.12)
Observations	148	148	148	148

Notes: Effect of income shocks on self-reported stress, cortisol levels, and positive and negative affect, OLS regressions with control variables. The dependent variables are self-reported stress, measured with a visual analog scale from 0–100; cortisol, measured in saliva in units of nmol/l; and positive and negative affect, measured with the PANAS scale. The estimates in each row represent an effect of interest (coefficient or combination of coefficients), and its robust standard error in parentheses. The estimates are obtained through parameter combinations of our main specification that identify the effect in question, as described in Section 2. The first row shows the difference between the "negative shock" and "always poor" groups; the second row between the "positive shocks" and "always rich" groups; and the third row shows the difference between these two effects. Control variables are included in all regressions and include family income, a dummy for being employed, and a dummy for being currently in debt. *** $p<0.001$, ** $p<0.01$, * $p<0.05$.

Table S6: Effect of income shocks on Stress, Cortisol Levels and Positive and Negative Affect (No Control Variables)

	Stress	Cortisol	Positive Affect	Negative Affect
Negative Shock vs.	−9.21	−59.96	−0.25	0.09
Always Poor	(5.43)	(39.50)	(0.17)	(0.10)
Positive Schock vs.	−8.16	−15.35	0.01	−0.11
Always Rich	(4.71)	(42.40)	(0.16)	(0.07)
Interaction	1.05	44.61	0.26	−0.20
	(7.26)	(57.95)	(0.23)	(0.13)
Observations	148	148	148	148

Notes: Effect of income shocks on self-reported stress, cortisol levels, and positive and negative affect, OLS regressions without control variables. The dependent variables are self-reported stress, measured with a visual analog scale from 0–100; cortisol, measured in saliva in units of nmol/l; and positive and negative affect, measured with the PANAS scale. The estimates in each row represent an effect of interest (coefficient or combination of coefficients), and its robust standard error in parentheses. The estimates are obtained through parameter combinations of our main specification that identify the effect in question, as described in Section 2. The first row shows the difference between the "negative shock" and "always poor" groups; the second row between the "positive shocks" and "always rich" groups; and the third row shows the difference between these two effects. Control variables are not included. *** $p<0.001$, ** $p<0.01$, * $p<0.05$.