Supporting Information Appendix for

Reducing Debt Improves Psychological Functioning and Changes Decision-Making in the Poor

Qiyan Ong*, Walter Theseira, Irene Y.H. Ng
*Corresponding author. Email: qyong@nus.edu.sg

This PDF file includes:

www.pnas.org/cgi/doi/10.1073/pnas.1810901116

Section 1. Debt relief program details

The researchers collaborated with Methodist Welfare Services, a charity in Singapore, to conduct the study. Although the charity is affiliated with the Methodist Church, to maintain tax-exempt status, charitable institutions in Singapore are generally required to serve the community at large, so program recipients were not restricted to any particular subgroup of low-income households in Singapore. The charity did not fund the research, and the charity did not exercise editorial control over the research.

Key details on the debt relief program are provided in the main text. This section addresses additional details which may be of interest.

Debt relief program design and implementation

The debt relief program was designed, funded, and implemented entirely by Methodist Welfare Services. The researchers did not provide any inputs on the program design or implementation. The charity has previously implemented programs for indebted clients, ranging from financial counselling to matching repayment schemes. For 2015, the charity implemented a generous one-off debt relief scheme, titled "the Getting-Out-of-Debt (GOOD) Programme", that would pay for up to SGD5000 in chronic debts, as part of Singapore's 50th anniversary celebrations. At 2015 purchasing power parity exchange rates, the debt relief program paid for the equivalent of USD5750 in debts. The charity carried out fundraising specifically to support the program. The researchers are not aware that any debt relief program of similar magnitude was implemented in Singapore in recent years, and there was no publicity significantly in advance of the program targeted at potential beneficiaries. Hence, we are confident beneficiaries are likely to regard the program as an unanticipated one-off event, and have little incentive to engage in strategic behavior, such as stockpiling debts.

Debt relief program application procedure

The actual size and distribution of debt relief was determined by the following procedure. Social workers were briefed on the debt relief program, and asked to recommend clients who had eligible chronic debts (owed for at least six months) for the program. Social workers contacted eligible clients, collected documentary proof of the eligible debts, and completed application forms on their behalf. Social workers were required to recommend the clearance of specific debts, following program guidelines. Charity program administrators then reviewed the application forms and had final discretion to approve and structure debt relief.

How do social workers decide which debts to recommend for relief? Social work in Singapore is a profession requiring an undergraduate social work degree or other recognized social work qualification for employment. Social work education, and on-the-job training, emphasizes a holistic client assessment framework as the basis for social work interventions. This assessment framework considers, amongst others, the needs, strengths and resources of the client, as well as personal factors such as motivation. The assessment framework is supported by standard texts, as well as by practice-oriented guides issued by the Singapore Ministry of Social and Family Development, the authority in charge of funding and supporting social services in Singapore (*1-2*). While social workers will naturally hold diverse views on how to best intervene with clients, social work training and the use of formal assessment frameworks provide some bounds on a social worker's discretion. Training and assessment frameworks make it unlikely that social workers would systematically grant debt relief only to those clients expected to "succeed" after debt relief. In fact, only a small minority of social workers in Singapore attribute poverty primarily to client behavior (3), so there are no cultural reasons to suppose that unobserved client potential is the determining factor in social worker decision-making. It is possible that social workers use discretion to vary the amount of relief given to otherwise similar clients – as discussed below, this discretion is a source of variation in our study. Our point is that there is little reason to suppose this discretion is systematically biased in favor of "successful" clients.

The professional judgment exercised by social workers provides idiosyncratic variation in the extent and amount of debt relief granted to clients. When a client owes more debts than can be cleared by the program, the social worker must decide which debts should be prioritized for relief. Or, there may be multiple debts within the program debt relief cap, but the social worker may believe that clearing some debts may undermine personal responsibility and/or encourage counterproductive behavior. There may also be debts excluded from debt relief, but which the social worker recommends nonetheless as exceptions to program policy (we exclude the small minority of beneficiaries with such debts from the analysis). There are likely differences of opinion between social workers on how to structure debt relief, and these differences of opinion are why variation in debt relief is idiosyncratic, and hence treated as a quasi-experimental source of variation in debt structure.

One concern is that some social workers may be more effective at obtaining benefits for clients than others; if so, more severely indebted applicants would seek assistance from those social workers, increasing the correlation between (endogenous) indebtedness and the extent and structure of debt relief. However, the social workers who completed debt relief applications were invariably the clients' existing assigned social worker, and it is unlikely clients could alter their social worker assignments to take advantage of the debt relief program. Clients in general cannot "shop around" for social assistance services, as national policy assigns clients to the nearest Family Service Centre (through which the debt relief program was administered). While exceptions do exist, given the unanticipated nature of the program and lack of advertising, it is unlikely that client sorting would have occurred specifically to take advantage of the debt relief program.

Application approval anticipation effects

Program applications were accepted from April 2015 onwards, while clients were notified by program administrators that their applications were successful in the last week of July 2015. Given the high application approval rate, a natural concern is anticipation effects: Did clients alter their behavior – in terms of consumption, borrowing, or psychological functioning – after application because they knew they would receive debt relief? In general, anticipation effects immediately after program application are unlikely for two reasons: clients had no basis for assuming their applications would be approved, and clients have limited ability to change consumption behavior to take advantage of pending debt relief.

First, the charity's program administrators did not know, at the time of program design and implementation, how many applications they would receive, or whether most applications would qualify for the stated program criteria. Program administrators budgeted for about 850 households, and implemented an initial program relief cap of SGD3000 per applicant. As applications were filed, program administrators determined there was no need to ration program benefits strictly, as only 656 applications were received by the close of applications. The program relief cap was subsequently raised to SGD5000 per applicant, given the excess available budget with 656 applicants (instead of 850). However, only charity program administrators knew this. Social workers and clients had no basis for assuming, at the time of application, that approval rates would be high. While it is possible that some social workers inferred approval rates would be high, it is unlikely a professional social worker would consider it ethical to advise clients to assume their application would be approved, since the client's trust and their professional relationship would be damaged if their advice turned out to be incorrect.

Second, most client debts are for bills in arrears, and not for consumer debt. Since creditors have no reason to believe the charity would soon repay the client's debts, it is improbable that creditors would have increased forbearance, extended more credit, or reduced collection activities simply because the debt relief application was pending. As clients knew that debt relief, if approved, would be paid directly to the creditor organizations rather than disbursed in cash, clients also had little reason to increase formal or informal cash borrowing in anticipation of receiving debt relief, since relief would not provide them with ability to repay such borrowing. Simply put, the nature and type of debt relief issued made it impractical for clients to anticipate debt relief, or to alter behavior in advance of receiving debt relief.

We note anticipation effects do exist on formal notification of debt relief approval, which we discuss in *(SI Appendix, Fig. S2)*. Also, we note that at the time of study recruitment, the researchers were informed in confidence by program administrators that approval rates were likely to be high. Based on this, we planned to interview every single applicant referred to us by social workers, and did not seek prior confirmation from program administrators on whether a particular client would be approved (In the end, 2 referred participants had their applications rejected). We avoided communicating this information on high approval rates to anyone other than the principal investigators and their research assistants. The surveys were conducted by contracted professional survey staff, who were never informed of (and have no reasons to care about) details of the debt relief program.

Section 2. Field study description

Overview

Our field study design captures participant data using a comprehensive household financial survey, augmented with measures of psychological functioning and economic decision-making, and administrative data provided by the charity. Study participants were provided with compensation for their time, and economic decision-making tasks were incentivized. The study's first wave was designed to capture data pre-debt relief, while the second wave captured data three months post-debt relief. The measures of psychological functioning and economic decision-making are described later in these supplementary materials.

Assistance with study recruitment

As discussed in the main paper, due to personal privacy regulations we could not solicit study participants from a "master list" of all applicants to the debt relief program. Therefore, with the assistance of the charity, we requested Family Service Centre social workers administering the debt relief program to ask applicants for consent to be referred to the researchers for the study. Social workers were told that study referral was voluntary and would not affect debt relief program eligibility or benefits. The 281 referrals to the study come from this process.

Sample selection and representativeness

The debt relief program received 656 applications, of which 619 were approved. Social workers referred 281 participants to the researchers, and we contacted and surveyed 238 participants in the first wave.

The main paper's analysis sample has 196 participants. We excluded 42 first-wave surveyed participants, for one or more of the following reasons: 22 participants received relief for a non-eligible debt type, indicating their debt circumstances were judged exceptional by social workers or program administrators; 18 participants were missing measurements in psychological functioning and/or or economic decision-making (participants could refuse to answer questions without penalty); 3 participants were lost to attrition between the first wave pre-debt relief survey and the second wave post-debt relief survey; 2 participants had their applications rejected. Out of the 42 excluded participants, 39 participants had exactly one of the above issues, and 3 participants had two of the issues.

To check for selection effects into the main paper's analysis sample, we rely on de-identified administrative data collected from the application forms for all applicants. *(SI Appendix, Table S1)* reports debt relief amount and structure, and basic demographic information, based on administrative data, by four sample groups: All Applicants, Approved Applicants, All Pre-Debt Relief Survey Participants, and Our Sample used in the main paper. Each successive group is a subset of the earlier. On observable characteristics, the final study sample is broadly similar to the population of all approved applicants: Aggregate debt relief amounts, debt accounts paid off in full, income, household size, and employment are all similar between the two groups.

Section 3. Variable construction and experimental procedures

Debt relief granted

The charity provided administrative data on debt relief applications, payments and demographic characteristics for all clients of the debt relief program. The data is de-identified for clients who are not study participants. The administrative data reports for each client, the structure and amount of debt relief (i) applied for, and (ii) paid to each creditor.

There are five categories of eligible debts: Housing arrears, Utilities arrears, Town Council arrears, Telco arrears and Hire-Purchase arrears. Considering the debt relief program as a whole, *(SI Appendix, Table S1)* shows the most common debt types paid were Utilities (50%), Housing (42%) and Telco (31%); the percentages in parentheses are the proportion of approved applicants receiving relief in those categories. Conditional on receiving debt relief in that category, the largest (average) debt relief grants were for the Hire-Purchase (SGD3120), Housing (SGD2377), and Telco (SGD1146) categories. As explained shortly, these figures reflect the total amount of debt relief granted within the category, and not the actual debt size owed to an individual creditor within the category.

(SI Appendix, Table S1) shows the final study sample is similar to the debt relief program as a whole in terms of the structure of debt relief granted, with the exception that the final study sample was slightly more likely to receive debt relief for Town Council and Telco debts, and did not receive any debt relief in Non-eligible debts, by design. The final study sample also received slightly larger debt relief grants for Housing debts and Town Council debts, and slightly smaller debt relief grants for Hire-Purchase debts. These minor differences suggest no substantive issues affect selection into the final study sample.

We measure the total amount of debt relief granted by aggregating all debt payments recorded in the administrative data for that client.

Debt accounts paid off

There is a distinction between debt accounts and debt relief categories. A debt account is defined as one specific account in arrears, owed to one creditor. The Telco and Hire-Purchase categories frequently contain multiple debt accounts, because there are several phone companies and consumer goods retailers in Singapore. Therefore, debt relief beneficiaries frequently owe debt accounts to more than one creditor within the Telco and Hire-Purchase categories, and debt account sizes are smaller than the total amount of debt relief granted in that category. The Housing, Town Council, and Utilities categories usually contain a unitary debt account (with minor exceptions). For Housing debt, the main creditor for debt relief beneficiaries is the Housing Development Board (HDB), the national public housing authority. While banks also issue mortgage loans, about two-thirds of all public housing mortgages are issued by the HDB, with lower-income households especially likely to have HDB loans due to strict underwriting standards for bank loans. Most debt relief applicants do not qualify for bank loans, and rely on the Housing Development Board for credit. For Town Council debt, the creditor is one of the local Town Councils in Singapore, which maintain and administer common areas in a public housing estate. A beneficiary may owe debts to more than one Town Council if they have moved between Town Council jurisdictions, but this is rare in our sample. Finally, for Utility debt, the creditor is SP Services, the national utility services billing agency. It is not normally possible for an individual to hold more than one household billing account with SP Services.

We use the debt accounts, rather than the debt category, as the measure of debt mental accounts in our study. This more closely matches the actual debt experience for participants, who receive billing statements from individual creditors separately, and who must pay bills or make repayment arrangements with creditors separately as well.

Measuring the number of debt accounts paid off is not straightforward because the administrative data records do not have a variable identifying when a debt account is completely paid off. The data released to us by the charity only identify payment amounts and the debt account to which payment was made. The original billing statements have been archived by the charity separately from the current administrative dataset because the statements contain personally identifying information. At present, no resources are available to transcribe the billing statements (which are hand-made scans of bills provided by applicants, and are difficult to transcribe by machine).

Nonetheless, we can still identify a debt account payoff with a high degree of accuracy, based on the program administrator's debt repayment policy. The program administrators reviewed each application individually before approval, and compared the debt relief requested with the original billing statements. Because the program had excess budget (due to receiving too few applications, as discussed earlier) the program relief cap was raised to SGD5000, and program administrators implemented a policy of making full repayment, based on submitted billing statements, for accounts which social workers only recommended for partial repayment. Recall that social workers had the discretion to allocate debt relief; program administrators simply "topped-up" allocated partial debt account payments into full debt account payments, provided such payments were still within the SGD5000 relief cap.

The administrative data records are consistent with this policy. The records show hardly any instances of rounding or regularities in payment, except for beneficiaries who receive debt relief at the program cap of SGD5000, and who have that relief paid entirely to one creditor. Otherwise, virtually every single payment is for a non-rounded value, down to the last few cents, consistent with payments being made to completely clear specific debt accounts based on billing statements.

Therefore, for the 75% of participants who receive debt relief of less than the final program cap of SGD5000, we consider any debt relief payment made to a specific creditor to imply full clearance of that debt account. For the 25% of participants receiving debt relief grants exactly at the cap of SGD5000, we assume that one debt account is only partially repaid. This means that beneficiaries receiving the capped debt relief amount, who only have one debt account paid, are treated as having zero debt accounts completely paid off; it is highly improbable any debt account would be worth exactly SGD5000. All the main results hold when we re-estimate our results excluding all beneficiaries receiving the capped debt relief amount of SGD5000 *(SI Appendix, Table S7).*

Cognitive functioning

Our measure of cognitive functioning is based on the Eriksen Flanker Task, adapted from (4-6), and deployed on a tablet. Prior to beginning, interviewers conducted a short briefing on the Flanker Task. The briefing text was also displayed on the tablet screen. Participants were informed they would see five arrows, displayed in a row, with each arrow pointing either left (“<”) or right (“>”). Participants were asked to determine the direction of the center arrow as accurately and as fast as possible, and to provide their answer by touching an on-screen button corresponding to the correct direction. *(SI Appendix, Fig. S4)* provides the screenshots of the task.

After completing the briefing, the participant was required to start the Flanker Task with 5 practice trials. Interviewers were instructed to guide participants through the practice trials, but to refrain from providing assistance during the actual trials. During the practice trials, feedback was provided after each trial to notify the participant whether their response was correct or incorrect. There is no correct/incorrect feedback for the actual trials. Participants proceeded directly to the 20 actual trials after completing the 5 practice trials. There is no requirement that participants complete all 5 practice trials correctly to proceed. The results from these 20 actual trials form the basis for our cognitive functioning measure.

Participants were allowed to skip the tablet-based tasks without penalty. Interviewers were also allowed to use discretion to skip tablet-based tasks for participants who appeared incapable of using the tablet or understanding the task. As mentioned in *(SI Appendix, Section 2)*, 18 participants out of 238 interviewed in the first wave had incomplete responses for at least some of the psychological functioning or economic decision-making tasks, and are excluded from the final study sample.

The following cognitive functioning measures are constructed from participants’ performance in the Flanker Task:

1. **Error rate:** The proportion of trials (out of 20 actual trials) for which the participant gave the wrong answer;
2. **Log of median reaction time:** Natural log of the median reaction time (out of 20 actual trials) at the participant level;
3. **Combined score:** A combination of the accuracy score and the reaction time score based on the calculation method reported in the NIH Toolbox (4-5). The combined score ranges from 0 to 10. A summary of the computation method follows:
    a. The combined score is composed of an accuracy vector and a reaction time vector. Each vector ranges from 0 to 5. The vectors are added to obtain the final combined score from 0 to 10. However, if the accuracy vector is 4 or less (representing an 80% accuracy rate; see below), then only the accuracy vector is used, and the reaction time vector is discarded. This limits affected participants to a combined score ranging from 0 to 5.
    b. The accuracy vector is computed by:
    $$0.125 * Number\ of\ Correct\ Trials$$
    All respondents aged 8 and older – our entire sample – are automatically given twenty correct trials. This is because the NIH Toolbox Flanker Task, whose computation method we adopt, has a combined test protocol for both children and adults. Children aged 3 to 7 are given a first set of twenty trials using fish stimuli, which are easier. Respondents aged 8 and above are assumed to score all fish stimuli trials correctly, and are automatically given twenty correct trials without the need to present the fish stimuli trials. Therefore, in our sample, the accuracy vector only ranges from 2.5 to 5. The 80% accuracy rate (see above) is calculated assuming the participant has 100% accuracy for the unpresented set of 20 fish stimuli trials. Hence, a participant must answer more than 12 out

of 20 Flanker trials correctly to have their reaction time vector added to their accuracy vector, for the combined score.

c. The reaction time vector is computed based on the participant's median reaction time for qualifying trials. A qualifying trial satisfies the following conditions: the trial must be correctly answered, must have a reaction time greater than or equal to 100ms, and must not take longer than 3 standard deviations more than the participant's mean reaction time. Although the NIH Toolbox method considers only incongruent trials as qualifying trials (4-5), the software we used did not record the congruent/incongruent state of the stimuli, and presented the congruent/incongruent stimuli in random order, so we could not recover congruency data. Therefore, we consider all trials as potentially qualifying. Since congruent trials are easier to answer correctly, this means our median reaction times may be faster when compared to other Flanker Task implementations. Finally, the NIH Toolbox method recommends only including participants with median reaction times between 100ms and 10,000ms in the sample (4-5). All participants in our study who completed the Flanker Task met this criterion. To produce a reaction time vector score bounded between 0 and 5, the participant-level median reaction times are censored with a minimum value of 500ms and a maximum of 3000ms, and then scaled between 0 and 5 using the formula in (4-5):

$$Reaction\,Time\,Score = 5 - \left(5 * \left[\frac{\log(RT) - \log(500)}{\log(3000) - \log(500)}\right]\right)$$

This produces a reaction time score where 5 indicates a fast median reaction time at or below the censoring value of 500ms, while 0 indicates a slow median reaction time at or above the censoring value of 3000ms. We were able to compute a valid reaction time score for all participants who completed the Flanker Task.

<u>Anxiety</u>

Our measure of negative affect is based on a battery of eight questions used in the DSM-IV criteria for Generalized Anxiety Disorder (GAD), described in (*7*). The battery of questions is administered via interview. Previous studies have established that GAD symptoms often arise in the poor, and that these symptoms are quantitatively distinct from the background stresses of poverty (*8*). This property makes GAD (symptoms) appropriate as a measure of psychological affect that may be impacted by scarcity.

(SI Appendix, Table S14) lists the eight GAD questions asked in the study. To be classified as exhibiting GAD symptoms, based on the DSM-IV criteria, the participant must reply "Yes" to the first two Questions, and must also reply "Yes" to at least 3 of the symptoms stated in the subsequent six Questions. The study variable GAD is a binary variable with the value 1 if the participant exhibits GAD symptoms, based on the DSM-IV criteria, and 0 otherwise.

<u>Risk aversion</u>

Our measure of risk aversion is the incentivized risk elicitation task designed by (*9*). The task requires participants to choose one of six lotteries listed in *(SI Appendix, Table S15)*, each with a 50-50 chance of winning the higher or the lower reward. Approximately 5 percent of participants were randomly selected to be paid based on their lottery choice. The selection probability for the risk choice task is independent from that of the intertemporal choice task.

Participants were told that a random draw would be conducted by the tablet at the end of the interview to determine if their incentivized task choices were selected for payment. However, the exact probability of being selected was not disclosed to participants or to the interviewers. Screenshots of the task are provided in *(SI Appendix, Fig. S5)*.

The six lotteries are ordered by decreasing intervals of *r*, the Constant Relative Risk Aversion parameter *(SI Appendix, Table S15)*. We chose this risk elicitation task, even though it is relatively simple (and yields less precise risk parameter estimates), because it is easily understood, especially by participants with poorer math skills, and is less likely to elicit random responses (*10*). We infer a participant's range of *r* from their observed choice. We also infer whether and to what extent *r* changes after debt relief from computing the difference between the CRRA parameter intervals of the participant's pre- and post-relief choices.

One issue with the CRRA parameter intervals is treatment of the upper and lower bounds of the extreme choices (option 1 and 6), which are positive and negative infinity in *r* respectively. While interval regression methods can fit models with intervals that are unbounded on one side, extremely high or low CRRA parameter values imply decision-making over risk that does not fit real-world behavior. Very high *r* suggests, for example, that the risk premium for investing should be extraordinarily large, while very negative *r* suggests that the decision maker is willing to pay well in excess of the expected value to buy even highly risky gambles.

Therefore, in the empirical analysis we implement upper and lower bounds of *r* for the CRRA parameter intervals for option 1 and 6. We implement an upper bound of $r = 10$, based on the maximum CRRA utility function curvature parameter assumed in general equilibrium model studies of the equity risk premium (*11*). Determining a reasonable lower bound of *r* is more difficult. There are few elicitations of negative *r* in the literature, perhaps because declining marginal utility of wealth (hence, risk aversion) appears more consistent with real world behavior. One widely used risk elicitation method that does elicit negative *r* is the multiple price list experiment in (*12*), whose lowest bounded interval of *r* is $-0.95 < r < -0.49$. Only 1% of choices, when made for real cash stakes, are in that range, with another 1% in the interval $r < -0.95$. Therefore, we implement a lower bound of $r = -1$, which implies – for example – that a decision maker would be indifferent between option 1's certain value of SGD28, and a gamble of SGD40 versus SGD0. All our analyses that use the CRRA parameter interval implement the $r = 10$ upper bound, and the $r = -1$ lower bound. We test the sensitivity of our results to more extreme lower bounds in *(SI Appendix, Table S16)*.

Present bias

Our measure of time discounting is an intertemporal choice task (also called a "money-earlier-or-later" task) using two incentivized multiple price lists (13-15), shown in *(SI Appendix, Table S17)*. Each price list has eight choices that require participants to trade-off receiving a varying smaller payoff earlier, versus a larger fixed payoff of SGD50 at a later date. The lowest payoff is SGD30. The first price list offers a choice between payoffs today versus one month later, while the second offers payoffs at six months versus seven months later.

Time discounting tasks may be affected by perceived differences in payment risk for later payouts. For example, participants may consider promises of later payment less credible, thereby choosing earlier payments to avoid uncertainty, and not because they have high discount rates (*16*). We took steps to ensure participants viewed the research study and promises of payment as credible. All participants were introduced to the study through their social workers, who obtained participant consent to have their contact information transferred to the research team. Besides satisfying personal data protection regulations, this step ensured participants knew their social workers were supportive of the study. When recruiting participants, the research team was careful to explain that the study was conducted by the National University of Singapore, the national flagship university. The researchers provided grocery vouchers to all participants as reimbursement for their time, demonstrating the study's financial ability to pay. The personal information sheet, which participants retained, contained the email and phone number of the researchers, as well as that of an independent National University of Singapore Institutional Review Board member. Finally, we attempted to minimize uncertainty over later payments. Participants who were entitled to a later payment were provided with a written claim form that listed the dollar amount they were entitled to, the payment issuance date, and contact information of the research team (again). The participants were promised that any payments would be delivered to them on the issue date to minimize inconvenience. All choice experiment payments were made in cash, minimizing any uncertainty over payment type, bank account availability, and so forth.

Approximately 5 percent of participants were randomly selected to be paid based on their choices; if selected, one of their sixteen price list choices was randomly chosen as the basis for payment. The selection probability for the intertemporal choice task is independent from that of the risk choice task. Participants were told that a random draw would be conducted by the tablet at the end of the interview to determine if their incentivized task choices were selected for payment. However, the exact probability of being selected was not disclosed to participants or to the interviewers. Participants who were selected for payment received cash payment on the spot, or at their homes, at the dates specified in the task. Screenshots of the task are provided in *(SI Appendix, Fig. S6)*.

A participant's choices in the two price lists allows us to infer their discount factor, and hence, whether they exhibit present bias. We assume when inferring discount factors that participants intend to consume payoffs immediately when received, and that consumption utility is locally linear over the payoff stakes (*15*). These assumptions are admittedly strong, and reflect the challenge of estimating "true" time discounting preference parameters from small-stakes cash experiments. Nonetheless, we believe the discount factor estimate is still useful if the change in time preferences between survey waves is driven by debt relief, rather than by other contemporaneous changes (e.g. calendar effects) to the participant's personal or financial situation. We address the question of calendar effects more thoroughly in *(SI Appendix, Table S8)*.

The discount factor is inferred from the largest earlier payoff the participant is willing to accept before they switch to the fixed later payoff. For example, if the participant accepts SGD49 today instead of SGD50 in one month, but then rejects SGD48 today instead of SGD50 in one month, we can infer their monthly discount factor is: (Largest Accepted Earlier Payment/Fixed Later Payment) = (49/50) = 0.98. Participants who always choose the later fixed payment of SGD50 are assumed to have a discount factor of 1, while participants who always choose the smaller earlier payoff are assumed to have a discount factor of (30/50) = 0.6.

A participant with a lower discount factor ("impatient") prefers smaller earlier payoffs to waiting for larger fixed payoffs. Present bias is characterized by dynamically inconsistent choices, that is, favoring smaller earlier payoffs *to a greater extent when the payoff is immediate,* compared to when the earlier payoff occurs in the future. Consider a participant who switches to the later payoff at SGD38 today instead of SGD50 in one month, and switches at SGD48 in 6 months instead of SGD50 in 7 months. This participant has a lower discount factor when the earlier payoff is immediate, compared to when the earlier payoff occurs at 6 months. This means the participant is more "impatient" the closer at hand the payoff is, and exhibits present bias.

Similar to (14), some participants in our study (24%) made inconsistent choices, switching multiple times between a smaller earlier payoff and the larger fixed payoff. We exclude these participants from the main analysis, but include these participants under alternative assumptions that identify their discount factors in *(SI Appendix, Table S4)*. Of the remaining participants, on average, 72% do not switch at all between the smaller payoff and the larger fixed payoff in any given price list. The high rate of non-switching between payoffs is not unusual. In (14), 40% of subjects did not switch between payoffs in any given price list. While this suggests a wider range of payoffs should be offered, the implied discount factor of choosing the smallest earlier payoff of SGD30 is already extremely low and indicates very high impatience.

In addition, discount factors can still be bounded even if the participant reports no switches within a price list. For example, a participant who always chooses a smaller earlier payoff, and never switches to the later larger fixed payoff of SGD50, should have a discount factor less than that implied by a switch at the lowest smaller payoff of SGD30. A participant who switches at the same point in both price lists, or who always chooses the larger fixed payment of SGD50, is classified as having no present bias. Because discount factors can only be broadly bounded for a significant proportion of participants, we restrict our main analysis to a binary indicator of whether the participant's choices reflect present bias. In *(SI Appendix, Table S5)*, we present an alternative analysis based on a continuous measure of the participant's present bias parameter.

Table S1. Debt relief program characteristics by sample group

Variables	Statistic	All Applicants (N=656)	Approved Applicants (N=619)	All Pre-Debt Relief Survey Participants* (N=238)	Our Sample** (N=196)
Debt Relief Amount (SGD)	Mean	2372	2488	2532	2548
	SD	1697	1669	1681	1671
Debt Relief Accounts Paid In Full	Mean	1.60	1.65	1.68	1.69
	SD	1.03	1.01	1.04	1.02
Proportion With No Household Income	%	16	14	11	11
Total Household Income (SGD)***	Mean	1817	1803	1860	1716
	SD	1150	1139	1295	1013
Household Size	Mean	4.32	4.36	4.44	4.41
	SD	1.81	1.79	1.75	1.75
Employed Applicants	%	62	62	60	60
Female Applicants	%	74	74	76	77
Proportion receiving debt relief and average debt relief granted in SGD (excluding relief = 0), by debt category****					
Housing	%	40	42	38	42
	Mean	2351	2377	2635	2602
	SD	1679	1686	1729	1750
Town Council	%	26	26	29	32
	Mean	661	677	749	796
	SD	857	866	954	992
Utilities	%	48	50	50	52
	Mean	922	935	925	938
	SD	1007	1011	996	1021
Telco	%	30	31	33	36
	Mean	1129	1146	1070	1134
	SD	941	944	985	1005
Hire Purchase	%	9	10	10	10
	Mean	3120	3120	2878	2983
	SD	1559	1559	1579	1671
Non-Eligible Debts	%	8	8	9	0
	Mean	2214	2260	2310	0
	SD	1752	1762	1876	0

* A small number of pre-debt relief surveys were conducted with referred clients whom, on further checks, were found to be non-applicants. These have been dropped from the data and are not reported here. Reported here are 2 referred clients whose applications were subsequently rejected.
** Demographic data for "Our Sample" in this table is based on administrative data to facilitate comparison across groups. Minor differences between the administrative data and our survey data exist for three reasons: (i) Several months have passed between debt relief application and the survey date; (ii) Administrative data reports dependants by group bins 1-2, 3-4 and 5 and above, which we replace with 1.5, 2.5 and 5.5 respectively to compute household size statistics in this table; (iii) we have not cleaned the administrative data in producing this table, to preserve the data in the format held by the charity.
*** Total household income (SGD) is conditional on positive income.
**** All Applicants includes 37 rejected applications. Some rejected applications have debts and proposed debt relief (which was not paid out) reported in administrative data. Other rejected applications have no administrative data on debt at all. For completeness, we include the available administrative data on debts and proposed debt relief for rejected applications in our summary statistics.

As discussed in *(SI Appendix, Section 2)*, there are four relevant sample groups: 656 applicants to the debt relief program; 619 applicants who received debt relief; 238 participants who completed initial interviews prior to receiving debt relief; and 196 participants in the analysis sample who received debt relief only for eligible debts, who answered all psychological functioning and economic decision-making questions, and who were contactable in the follow-up wave. We have demographic and debt relief administrative data for all groups, but we only have survey data measures for the 238 participants who completed our initial interview. The table shows that our participants have similar demographics, and received similar debt relief benefits, as non-participant clients. Non-eligible debt relief is zero for our study participants, as we excluded participants who received relief for such debts.

Table S2. Main analysis excluding high-debt participants

	Error Rate (Flanker)	Response Time (Flanker)	Combined Score (Flanker)	GAD	Risk Aversion	Present Bias
Debt Relief Amount	-0.030***	-0.043**	0.234***	-0.016	-0.076	0.025
	(0.006)	(0.019)	(0.054)	(0.018)	(0.081)	(0.019)
Debt Accts Paid Off	-0.032***	-0.094***	0.336***	-0.116***	-0.247*	-0.097***
	(0.010)	(0.025)	(0.088)	(0.025)	(0.145)	(0.032)
Constant	0.174***	0.450***	6.318***	0.768***		0.436***
	(0.008)	(0.022)	(0.066)	(0.018)		(0.027)
Observations	372	372	372	372	186	280
No. of Participants	186	186	186	186	186	140

See main paper Table 2; regression models are identical.

A small number of participants have debt levels significantly exceeding the sample average. The 95th percentile by debt size is SGD22450, while the average and median debt is SGD6257 and SGD3574 respectively. These large debts are generally from past due mortgages, and may complicate our analysis as they may be correlated with unobserved confounding factors.

Mortgage lending policies applied by the Housing Development Board (the public housing authority, and the main creditor for housing debts in our sample) and by financial institutions in Singapore are generally conservative. Current lending policy limits mortgage repayments to 30% of a borrower's gross monthly income. Given the low household incomes in our study population, the presence of such large mortgage debts suggests the household's income has sharply decreased at some point prior to the survey period, and/or that the household experienced significant life events such as divorce, death, or disability. Large mortgage debt also suggests forbearance on foreclosure, which is more likely when the household presents extenuating life events to the lender.

These factors raise the concern that such high-debt participants may not be representative of low-income households in general, and/or are likely to have unobserved confounding factors that significantly affect indebtedness, psychological functioning, and decision-making. As a robustness check, we repeat the main analysis after excluding participants with initial debt value above the 95th percentile. The results are broadly consistent with the main analysis on the unrestricted sample, and show that the effects are not driven by participants with extraordinarily large debts.

Table S3. Psychological functioning and decision-making measures, pre and post debt relief

	Pre Debt Relief	Post Debt Relief
Error Rate (%)	17.321	4.056
	(1.428)	(0.691)
Median Response Time (s)	1.860	1.266
	(0.082)	(0.450)
Combined Score	6.290	7.545
	(0.129)	(0.842)
GAD (%)	78.06	52.55
	(2.964)	(3.576)
No-risk lottery of (28,28) (%)	32.14	20.408
	(3.344)	(2.886)
High-risk lottery of (72,2) (%)	16.837	29.082
	(2.680)	(3.252)
Present Bias (%)	43.624	32.886
	(4.076)	(3.862)

196 participants except where noted. Mean values reported. Standard error of the mean in parentheses. Present Bias is based on 149 participants with consistent time discounting responses.

Table S4. Present bias analysis including participants with inconsistent choices in the time discounting experiment

	Present Bias		
	Only Consistent Participants	Lowest Switching Point, Full Sample	Highest Switching Point, Full Sample
Debt Relief Amount	0.029	0.013	0.006
	(0.017)	(0.017)	(0.016)
Debt Accounts Paid Off	-0.100***	-0.090***	-0.051*
	(0.032)	(0.027)	(0.030)
Constant	0.430***	0.462***	0.415***
	(0.026)	(0.022)	(0.022)
Observations	298	392	392
Number of Participants	149	196	196

See main paper Table 2; regression models are identical.

A participant's discount factor cannot be bounded precisely when the participant is inconsistent and switches multiple times between the earlier and the later payment in the same multiple price list. While the main results use only participants with precisely bounded discount factors (switching once, or never, within the same price list), the robustness checks here include inconsistent participants under assumptions that bound a discount factor for such participants.

The left-most column replicates the Table 2 "Present Bias" column in the main paper, which reports the effect of debt relief amount and debt accounts paid off, on changes in the participant's present bias. The left-most column thus uses only consistent participants. The middle column uses all participants, and assumes that the lowest recorded switching point within a price list represents the participant's true intention. The right-most column uses all participants and assumes that the highest recorded switching point within a price list represents the participant's true intention. The different sample restrictions show broadly similar results, although the relationship between debt accounts paid off and present bias is weaker and significant only at the 10% level when the highest switching point assumption is used. However, note that the highest switching point assumption generates high discount factors (and hence a low likelihood of categorization as having present bias), because many inconsistent participants start with choosing the $50 later payment and then switch back to smaller earlier payments.

Table S5. Present bias analysis using the (β,δ) quasi-hyperbolic model parameters

Panel A: Summary Statistics of δ and β from Participant Choices in the Time Discounting Experiment

	Pre Debt Relief	Post Debt Relief	T-Statistic, P-Value (two-tailed)
Consistent Sample (n=149)			
δ	0.896	0.873	t=1.204, p<0.230
β	0.944	0.956	t=-0.545, p<0.586
Proportion with Present Bias (β<1)	44%	33%	z=1.907, p<0.057
Full Sample, Lowest Switching Point (n=196)			
δ	0.883	0.864	t=1.083, p<0.280
β	0.946	0.958	t=-0.623, p<0.534
Proportion with Present Bias (β<1)	46%	34%	z=2.147, p<0.014
Full Sample, Highest Switching Point (n=196)			
δ	0.895	0.885	t=0.629, p<0.530
β	0.965	0.961	t=0.205, p<0.838
Proportion with Present Bias (β<1)	42%	34%	z=1.769, p<0.077

Panel B: The Effect of Changes in Debt Structure on the β Present Bias Measure

	Only Consistent Participants	Lowest Switching Point, Full Sample	Highest Switching Point, Full Sample
Debt Relief Amount	-0.007	-0.008	-0.005
	(0.008)	(0.007)	(0.008)
Debt Accounts Paid Off	0.015	0.018*	0.005
	(0.011)	(0.010)	(0.011)
Constant	0.946***	0.947***	0.965***
	(0.011)	(0.009)	(0.010)
Observations	298	392	392
Number of Participants	149	196	196

See main paper Table 2; regression models are identical.

To analyze continuous measures of participant time discounting preferences, we adopt the (β,δ) quasi-hyperbolic model proposed by (*17*). In the (β,δ) model, β is an additional discount factor applied to any payoff received later than the present, while δ is the standard discount factor over the payoff delay period (one month in our time discounting money-earlier-or-later experiments). Therefore, indifference between (C_0), the earlier payoff today and (C_1), the payoff one month later is given by:

$U(C_0) = \beta\delta U(C_1)$

While indifference between (C_6) the payoff at six months and (C_7) the later payoff at seven months is given by:

$\beta U(C_6) = \beta\delta U(C_7)$

Assuming utility is locally linear, the one-month discount factor δ is obtained from the ratio of the earlier to the later payment at the point of indifference, based on the six vs. seven month choice (so that β cancels out): $C_6/C_7 = \delta_{6,7}$. β is obtained from $(C_0/C_1)/(C_6/C_7) = \beta\delta_{0,1}/\delta_{6,7}$. As δ is assumed to be constant over any one-month delay, $\beta\delta_{0,1}/\delta_{6,7}$ simply reduces to β (14). A participant with present bias has $0 < \beta < 1$.

Panel A summarizes changes in β and δ by the three different time discounting samples *(See SI Appendix, Table S4)*. Overall, there is little change in the average one-month discount factor δ, or in the average β parameter. However, the proportion of participants who exhibit present bias (have a β parameter between 0 and 1) is significantly reduced after debt relief, in all three samples.

Panel B uses the regression model from the "Present Bias" column in Table 2 of the main paper, regressed on the β parameter. As higher values of β indicate a reduction in present bias, the results show a greater number of debt accounts cleared is associated with a reduction in present bias. The results are not statistically significant for the main analysis sample which uses only consistent participants, but they are of similar magnitude to the results using the full sample based on the lowest switching point, which is significant at the 10% level.

Table S6. Main analysis based on actual debt structure change instrumented with debt relief

	Cognitive Functioning (Combined Score)		GAD		Risk Aversion		Present Bias	
	FE OLS	FE IV	FE OLS	FE IV	INTERVAL OLS	INTERVAL IV	FE OLS	FE IV
Actual Debt Size	-0.286***	-0.201	0.023*	-0.010	0.136*	0.087	-0.017	-0.070**
	(0.058)	(0.123)	(0.013)	(0.036)	(0.070)	(0.142)	(0.020)	(0.034)
Actual Debt Accts	-0.123	-0.732***	0.123***	0.254***	0.188	0.464	0.072*	0.241***
Outstanding	(0.099)	(0.266)	(0.024)	(0.064)	(0.141)	(0.345)	(0.038)	(0.084)
Constant	8.758***		0.195***				0.275**	
	(0.269)		(0.075)				(0.108)	
Observations	392	392	392	392	196	196	298	298
No. of Participants	196	196	196	196	196	196	149	149

See main paper Table 2. FE IV is estimated using a 2SLS model with actual debt size and actual debt accounts outstanding as the endogenous variables, instrumented by the debt relief amount and debt relief accounts paid off. The FE IV models use STATA code developed by (18) while the Interval IV model uses STATA code developed by (19). CRRA is estimated using conditional mixed process model and the independent variables used are the change in actual debt size and change in actual debt accounts outstanding. Actual debt size and debt relief amount are in SGD Thousands. Robust standard errors are in parentheses. *p<0.10, **p<0.05, ***p<0.01

The main analysis' OLS results using the amount and structure of debt relief granted provide the average effect of debt relief intensity on the outcomes of interest. Conceptually, this is similar to an intent-to-treat analysis, where the final "treatment" status – changes in actual debt by the three-month mark, post-debt relief – is excluded from the analysis. The main analysis is based on these results as we believe they are of the most general interest, since policymakers can control the structure and implementation of debt relief, but have less ability to ensure that relief-induced changes in debt structure will persist over time.

The intent-to-treat estimates in the main analysis differ from the effect of actual changes in debt structure for two reasons. First, measurement error exists, because changes in actual debt structure are collected via our survey, whereas debt relief data is based on administrative records that document all payments made by the charity under the debt relief program. Moreover, the pre-debt relief survey is conducted right before debt relief is granted, while debt relief applications are based on the participant's debt position a few months prior. Second, participants may use the resources provided by debt relief to borrow more, or to restructure their debts further. These choices are likely related to changes in psychological functioning or decision-making, confounding the analysis further.

An OLS model estimating the effect of changes in actual debt structure will be biased because the change in actual debts is endogenous. To address endogeneity, we use an instrumental variables model, where the amount of debt relief, and number of accounts paid off, is used as an instrument that exogenously alters the actual change in the participant's debt amount and debt accounts. Debt relief grants are an ideal instrument for changes in actual debt because (i) debt relief grants strongly predict changes in actual debt structure, pre- and post-debt relief, and (ii) debt relief grants by themselves do not affect outcomes, except through the fact that they alter actual debt structure.

The IV model estimates the average causal effect of changes in debt structure for beneficiaries whose actual debt structure was most strongly influenced by the debt relief grant (*20*). Note that the coefficient signs are expected to reverse when the independent variable is the change in actual debt, compared to the main analysis' independent variable, which is the amount of debt relief granted: This is because a SGD1000 debt relief grant is equivalent to a *minus* SGD1000 change in actual debt.

The IV results are generally consistent with the main analysis' intent-to-treat estimates, and show that reduction in debt accounts, but not in debt amounts, cause improvements in psychological functioning. The IV model's effect of debt accounts is more strongly negative, compared to either the OLS using changes in actual debt, or the main analysis' intent-to-treat estimates. This is consistent with bias towards zero affecting OLS estimates of debt structure on psychological functioning. Such bias might result if those with more debt accounts also have greater unobserved ability to manage debts, which prevents psychological functioning from deteriorating under the mental load of multiple debt accounts. While there is still no evidence that debt structure explains risk aversion, reduction in debt accounts significantly reduces present bias, and as with the effects on psychological functioning, the OLS estimates appear biased towards zero. Although the IV results on present bias do suggest debt amount offsets the effect of debt accounts, the magnitude of debt amount effects is small: a change of ~SGD3400 in debt amount offsets the effect of just one debt account on present bias.

Table S7. Main analysis excluding participants at the relief cap

	Error Rate (Flanker)	Response Time (Flanker)	Combined Score (Flanker)	GAD	Risk Aversion	Present Bias
Debt Relief Amount	-0.008	-0.017	0.134	-0.018	0.107	0.018
	(0.013)	(0.041)	(0.117)	(0.033)	(0.250)	(0.047)
Debt Accts Paid Off	-0.048***	-0.127***	0.421***	-0.112***	-0.420*	-0.099**
	(0.014)	(0.033)	(0.119)	(0.031)	(0.233)	(0.044)
Constant	0.160***	0.454***	6.380***	0.771***		0.465***
	(0.008)	(0.024)	(0.075)	(0.019)		(0.029)
Observations	294	294	294	294	147	226
No. of Participants	147	147	147	147	147	113

See main paper Table 2; regression models are identical.

As discussed in *(SI Appendix, Section 3)*, the program administrators implemented a policy of completely paying off all eligible debt accounts recommended in each application for relief. The exception is applications whose total eligible debts exceed the program relief cap of SGD5000. In these cases, program administrators paid off recommended debt accounts, where possible, and assigned the remainder of relief – up to the SGD5000 cap – to one of the remaining debt accounts. We assume for relief-capped participants that one of their debt accounts is only partially repaid, and we do not include this partially repaid account in our measure of debt accounts paid off.

The number of debt accounts paid off for relief-capped participants may be measured with error. We cannot check the amounts paid off against actual billing statement records, which are retained by program administrators in their internal archives. Therefore, as a robustness check, this Table reports the main analysis results, restricting the sample to exclude the 25% of participants who received debt relief at the program cap of SGD5000.

The results are broadly similar to that of the main analysis, in terms of the effect of debt accounts paid off. However, there is no longer any significant effect of debt relief amount on any psychological functioning measure. This is likely because the excluded relief-capped participants are often measured with zero debt accounts paid off. For these excluded relief-capped participants, all benefits of receiving debt relief are captured in the coefficient on debt relief amount. Excluding relief-capped participants from the sample thus has the same effect as restricting the analysis to participants with at least one debt accounts paid off. In this restricted sample, debt accounts paid off entirely drive the results.

Table S8. Main analysis by calendar effects subgroups

	Error Rate (Flanker)	Response Time (Flanker)	Combined Score (Flanker)	GAD	Risk Aversion	Present Bias
Panel A: Participants Interviewed Outside Festivals and Holidays Only						
Debt Relief Amount	-0.022***	-0.030	0.160***	-0.013	-0.253***	0.018
	(0.006)	(0.019)	(0.058)	(0.024)	(0.087)	(0.019)
Debt Accts Paid Off	-0.023**	-0.088***	0.308***	-0.126***	0.054	-0.085**
	(0.010)	(0.027)	(0.092)	(0.029)	(0.169)	(0.038)
Constant	0.136***	0.400***	6.611***	0.826***		0.410***
	(0.008)	(0.022)	(0.074)	(0.022)		(0.031)
Observations	268	268	268	268	134	208
No. of Participants	134	134	134	134	134	104
Panel B: Muslim Participants						
	Error Rate (Flanker)	Response Time (Flanker)	Combined Score (Flanker)	GAD	Risk Aversion	Present Bias
Debt Relief Amount	-0.022***	-0.022	0.168***	-0.013	-0.216**	0.031
	(0.007)	(0.018)	(0.060)	(0.023)	(0.086)	(0.026)
Debt Accts Paid Off	-0.029**	-0.087***	0.264**	-0.114***	-0.157	-0.057
	(0.012)	(0.030)	(0.108)	(0.030)	(0.170)	(0.042)
Constant	0.142***	0.377***	6.651***	0.817***		0.412***
	(0.009)	(0.025)	(0.082)	(0.023)		(0.037)
Observations	232	232	232	232	116	166
No. of Participants	116	116	116	116	116	83
Panel C: Non-Muslim Participants						
	Error Rate (Flanker)	Response Time (Flanker)	Combined Score (Flanker)	GAD	Risk Aversion	Present Bias
Debt Relief Amount	-0.032***	-0.073**	0.294***	-0.030	0.037	0.027
	(0.010)	(0.030)	(0.075)	(0.028)	(0.120)	(0.023)
Debt Accts Paid Off	-0.045**	-0.108***	0.496***	-0.110***	-0.318	-0.158***
	(0.018)	(0.038)	(0.139)	(0.041)	(0.235)	(0.041)
Constant	0.210***	0.543***	5.884***	0.710***		0.452***
	(0.013)	(0.035)	(0.100)	(0.028)		(0.032)
Observations	160	160	160	160	80	132
No. of Participants	80	80	80	80	80	66

See main paper Table 2; regression models are identical.

The evidence from our pre-post debt relief comparison is potentially confounded by calendar effects associated with seasons or festivals. We cannot address this through time-based fixed effects in the regression model, because all participants experienced debt relief at the same time. However, we are able to check the robustness of our results to calendar effects by subgroup analysis of (i) participants who were interviewed outside of festival/holiday periods, and (ii) Muslims/non-Muslims, due to the potential influence of the Ramadan religious holiday, which is the fasting month for Muslims.

Before we begin the subgroup analysis, we discuss why confounds due to calendar effects are expected to be minimal even in the full study sample. Singapore is a multi-ethnic, secular, urban country. While there are several holidays and festivals throughout the year, there are no prolonged festivities where patterns of daily life are altered, as may be the case in mono-cultural or agrarian settings. Our participants are multi-ethnic and multi-religious, limiting the effects of any particular festival or holiday on our sample. We designed the fieldwork to further minimize calendar effects, as discussed below. There is little variation in weather year-round, except for the biannual monsoon which increases rainfall. The effects of weather are not considered as there is no significant agricultural sector, food is imported from multiple sources ameliorating the effect of weather shocks on food prices, and, homes, offices and shops are well-adapted to the constant heat, humidity and rainfall.

Potential calendar effects from festivals and holidays

The study fieldwork period spanned from 11 July 2015 through 6 January 2016. The pre-debt relief survey fieldwork took place from 11 July 2015 through 15 August 2015, while the post-debt relief fieldwork took place from 31 October 2015 through 6 January 2016 (79% of cases were completed by 1 Dec 2015).

Singapore distributes public holidays to cover key observances for each major ethnic and religious group. During this period, there were six public holidays: Hari Raya Puasa, 17 July 2015 (Muslim religious festival); National Day, 9 August 2015 (secular public holiday); Hari Raya Haji, 24 September 2015 (Muslim religious festival); Deepavali, 10 November 2015 (Hindu religious festival); 25 December 2015 (Christian religious festival); and, New Year's Day, 1 January 2016 (secular public holiday).

Most of our participants (95%) are from the three major ethnic groups in Singapore: Chinese, Malays, and Indians. The Chinese are generally neither Muslim nor Hindu, and there were no Chinese-related public holidays during this period. Malays are predominantly Muslim. Indians are Hindu, Muslim, or hold other religious beliefs. Christians are in the minority in our sample, and in Singapore.

Calendar-based exclusion analysis

We designed the fieldwork to minimize calendar effects by instructing surveyors to avoid interviewing subjects during the six public holidays noted above. Nonetheless, to address any exceptions to the fieldwork protocol, and for additional robustness, Panel A excludes all participants interviewed on either a public holiday or the eve thereof. In addition, for the Hari Raya Puasa and Christmas-to-New Year's Day holidays we exclude additional days.

Hari Raya Puasa marks the end of the Muslim month of Ramadan where Muslims fast during daylight hours. While we were concerned fasting might affect our psychological and decision-making measures, we aimed to complete fieldwork before early August, when the charity was scheduled to start issuing debt relief approval notices, and debt relief payments. Hence, we started fieldwork on 11 July; starting fieldwork after Hari Raya Puasa on 17 July would have reduced pre-August fieldwork to only two weeks. To minimize potential confounding we instructed surveyors to prioritize interviewing non-Muslim participants during Ramadan, and if needed, to interview Muslim participants in the evening after the fast ends. As Singapore is a secular, Muslim-minority state, there are no publicly observed restrictions on work hours, shopping hours, or general behavior during Ramadan. Nonetheless, to address possible unobserved spillover effects, in the calendar-based exclusion specification we exclude all participants interviewed during Ramadan, from the start of fieldwork on 11 July 2015, and up to and including Hari Raya Puasa on 17 July 2015.

Singapore is also a secular, Christian-minority state. However, the end of the year, including Christmas, is a secular holiday period where households take vacation time and engage in festive activities. To minimize potential confounding effects we instructed surveyors to avoid interviewing participants during the Christmas to New Year's Day period. To address possible exceptions to the fieldwork protocol, and for additional robustness, we exclude all participants surveyed starting one week prior to Christmas, from 18 December 2015 onwards, and up to and including New Year's Day on 1 January 2016.

In total, these calendar-based criteria exclude 62 participants from the study sample, leaving 134 participants in Panel A. The majority of excluded cases – 39 participants – are non-Muslims. This is in part due to our fieldwork protocol, which prioritized interviewing non-Muslim participants during Ramadan. The effects of debt relief are broadly comparable to that found in the main study sample. The coefficient estimates are smaller (closer to zero) for the cognitive functioning outcome measures, although they remain statistically significant in all cases; statistical significance falls to the 10% level for the effect of debt relief amount on log median response time. The effects on GAD and on present bias are similar. However, the effects are notably different for risk aversion, where larger debt relief amounts now reduce measured risk aversion. This change in the effect on risk aversion is driven by differences between Muslim/non-Muslim participants, and is discussed more below.

<u>Muslim/Non-Muslim subgroup analysis</u>

To address any remaining concerns that fasting during Ramadan may affect our results, we spilt the sample into the 116 Muslim and 80 non-Muslim participants, and estimated all models separately on each subgroup. The effects of changes in debt relief are similar by Muslim/non-Muslim subgroup, although the magnitude of the coefficient estimates are generally larger, and at times more statistically significant, for non-Muslim participants. The only notable difference is that larger debt relief amounts reduce measured risk aversion in the Muslim subgroup. This is due to the Muslim subgroup having greater risk aversion, pre-debt relief. As a result, the Muslim subgroup exhibits larger average reductions in risk aversion, as post-debt relief risk aversion is similar between the Muslim/non-Muslim subgroups. Further studies on heterogeneity between groups in risk attitudes, and in other decision-making measures, may be useful. It is beyond the scope of this paper to examine the reasons for such group-based heterogeneity in risk aversion.

Table S9. Main analysis excluding participants receiving full debt relief

	Error Rate (Flanker)	Response Time (Flanker)	Combined Score (Flanker)	GAD	Risk Aversion	Present Bias
Debt Relief Amount	-0.019***	-0.030*	0.148***	-0.015	-0.152*	0.031
	(0.006)	(0.017)	(0.048)	(0.020)	(0.080)	(0.020)
Debt Accts Paid Off	-0.045***	-0.084***	0.398***	-0.115***	-0.208	-0.096**
	(0.010)	(0.028)	(0.096)	(0.028)	(0.164)	(0.040)
Constant	0.159***	0.398***	6.453***	0.817***		0.434***
	(0.008)	(0.025)	(0.073)	(0.020)		(0.031)
Observations	284	284	284	284	142	216
No. of Participants	142	142	142	142	142	108

See main paper Table 2; regression models are identical.

We address many of the endogeneity issues that plague cross-sectional analyses of the effect of debt on psychological functioning and decision-making by exploiting the quasi-experimental variation generated by social worker decisions on debt relief. However, a small number of participants had virtually all eligible debts cleared, resulting in a high correlation for these participants between debt relief granted and their initial debt structure. This raises the concern that estimates of the effects of debt relief will be affected by endogeneity.

To check the robustness of our results to this source of endogeneity, we repeat the main analysis on a restricted sample that excludes 54 participants (28% of the final study sample) who received practically full debt relief (more than 90% of eligible debts cleared), and thus are the most affected by endogeneity from initial debt structure. The results based on the restricted sample are broadly similar to the main analysis on the full sample, and suggest that our estimates are not significantly affected by endogeneity transmitted through the correlation between debt relief granted and initial debt structure.

Table S10. Liquidity constraints: Analysis of homeowners receiving housing debt relief

Panel A: Proportion of Homeowners with Present Bias, by Debt Relief Status and by Housing Relief (N=84)

	Debt Relief in Housing Arrears	No Debt Relief in Housing Arrears	[Housing Relief] - [No Housing Relief] z-statistic, p-value (two-tailed)
Pre Debt Relief	46%	60%	z = 1.242, p < 0.214
Post Debt Relief	41%	32%	z = -0.820, p < 0.413
[Pre-Relief] - [Post-Relief] z-statistic, p-value (two-tailed)	z = 0.469, p < 0.639	z = 2.692, p < 0.007	

Panel B: Effect of Changes in Debt Structure on Present Bias, in Homeowners by Housing Relief

	Debt relief in housing arrears	No debt relief in housing arrears
Debt Relief Amount	0.035	-0.007
	(0.026)	(0.039)
Debt Accounts Paid Off	-0.143	-0.148**
	(0.092)	(0.056)
Constant	0.452***	0.598***
	(0.055)	(0.044)
Observations	74	94
Number of Participants	37	47

See main paper Table 2; regression models are identical.

Changes in housing equity may result in significant improvements in liquidity constraints, which are a confounding factor. *(SI Appendix, Table S1)* indicates that 42% of our analysis sample received an average of SGD2602 in relief for housing arrears. Housing relief effectively increases housing equity among homeowners who owe arrears for past due mortgage payments.

The effects of housing equity on liquidity from secured credit are limited in our context. As of 2015, an estimated 79% of the Singapore resident population lives in flats sold by the Housing Development Board ("HDB flats"), Singapore's public housing authority (*21*). Except for mortgage loans, financial institutions in Singapore are restricted by policy from issuing secured credit against housing equity in HDB flats. All the homeowners in the sample own HDB flats only, and thus cannot access home equity credit from financial institutions, regardless of debt relief.

Nonetheless, homeowners might borrow from informal lenders against housing equity, plan to sell their HDB flats to realize liquidity, or otherwise behave as though liquidity constraints were eased. Hence, we test for whether there are systematic differences in present bias by housing relief status amongst homeowners in the sample.

The full analysis sample of 196 participants includes 108 homeowners (who owned homes in both waves); 48 homeowners received housing relief while 60 did not. Restricting the analysis to homeowners with consistent choices in the money-earlier-or-later experiment reduces the sample to 84 homeowners, of whom 37 received housing relief while 47 did not.

Panel A shows that the proportion of homeowners exhibiting present bias does not differ significantly by housing relief status, both pre- and post-debt relief. While both housing relief and non-housing-relief groups exhibited reductions in present bias post-debt relief, these reductions were only statistically significant in the non-housing-relief group (z = 2.692, p < 0.007). This suggests changes in housing equity-related liquidity are unlikely to explain our results, since there is no evidence that reductions in present bias are larger among housing-relief homeowners.

Panel B reports the regression analysis of changes in debt structure on present bias, in the homeowner subsample. Although the estimates are not precise, due to the small subsample, the point estimates of the effect of debt account clearance on present bias are similar between housing-relief homeowners and non-housing-relief homeowners. The effects of debt account clearance are statistically significant only for non-housing-relief homeowners. The effects of debt relief amount are small and insignificant, across housing relief status. Taken together, this suggests the effects of changes in debt structure on present bias that we document in the main analysis are not explained by changes in housing equity-related liquidity.

Table S11. Impact of psychological improvements on risk aversion and present bias

	Risk Aversion		Present Bias	
	(1)	(2)	(3)	(4)
Cognitive Functioning	-0.109	-0.026	-0.025	-0.000
(Combined Score)	(0.103)	(0.126)	(0.022)	(0.028)
Loses GAD	-1.209***	-0.901*	-0.015	0.103
	(0.422)	(0.468)	(0.102)	(0.116)
Debt Relief Amount		-0.072		0.025
		(0.079)		(0.019)
Debt Accounts Paid Off		-0.116		-0.111***
		(0.152)		(0.038)
Constant			0.561***	0.430**
			(0.151)	(0.177)
Observations	196	196	298	298
Number of Participants	196	196	149	149

See main paper Table 2.

Laboratory experiments show exogenous shocks to fear, anxiety, and cognitive ability increase risk aversion and present bias, while observational studies find lower cognitive ability is linked with greater risk aversion and present bias (*22*-*26*). Simultaneously testing the role of cognition and psychological affect on decision-making is important, in line with dual process theories proposing decisions result from the combination of the patient, risk neutral deliberative system, as well as the impulsive, risk-averse emotional system (*24*).

However, a causal test of these hypotheses is challenging, because (i) a participant's psychological functioning and personal finances change contemporaneously after debt relief; (ii) both types of changes may independently affect decision-making; and, (iii) psychological functioning and personal finances may be jointly determined during the post debt relief period. In particular, (ii) means debt relief fails the exclusion restriction as an instrumental variable for estimating the causal effect of psychological functioning, since debt relief independently affects decision-making through altering the participant's personal finances. We lack sufficient variables to instrument for both psychological functioning and personal finances. Therefore, the analysis here is descriptive rather than causal.

We run OLS regression models where changes in the CRRA parameter interval, and in present bias, are regressed on the debt relief induced changes in psychological functioning and changes in debt relief structure. We measure the change in cognitive functioning with the Flanker Task combined score, and we measure whether the participant stops exhibiting GAD symptoms after debt relief with an indicator variable. We do not control separately for any effect of starting to exhibit GAD symptoms post-debt relief, as only 8 participants did so.

The evidence is consistent with poverty affecting risk aversion through poverty-induced changes in negative affect. Participants with reductions in GAD symptoms exhibited lower risk aversion post-debt relief, even after controlling for the structure of debt relief and changes in cognitive functioning. There is no evidence that poverty-induced changes in executive function (as measured by the Flanker Task) are related to changes in either risk aversion or present bias. Other aspects of cognitive functioning – which we were unable to measure – may yet affect decision-making.

Table S12. Consistency of choices in the time discounting experiment

	Consistent Responses in the Time Discounting Experiment (1: Yes, 0: No)			
	(1)	(2)	(3)	(4)
Combined Score (Flanker)	0.029**		0.025*	
	(0.014)		(0.014)	
GAD		-0.091	-0.069	
		(0.056)	(0.057)	
Debt Relief Amount				0.017
				(0.013)
Debt Accounts Paid Off				0.012
				(0.021)
Constant	0.667***	0.924***	0.740***	0.834***
	(0.095)	(0.036)	(0.114)	(0.016)
Observations	392	392	392	392
Number of Participants	196	196	196	196

See main paper Table 2.

Decision makers with greater cognitive ability make more consistent choices in economic experiments (*24-25*, *27*). Therefore, debt relief induced improvements in cognitive functioning may improve consistency of choices in the time discounting money-earlier-or-later experiment. In that experiment, a decision maker chooses between a smaller, earlier payment that varies between SGD49 to SGD30, versus a fixed later payment of SGD50. As *(SI Appendix, Table S17)* indicates, the earlier payment becomes smaller as the participant progresses through the choice list. A participant who switches from the smaller earlier payment to the fixed later payment should, from that point on, always prefer the fixed later payment. For example, it is inconsistent to choose, SGD50 later over SGD34 earlier, but then to choose SGD30 earlier instead of SGD50 later. Multiple switches indicate choice inconsistency.

First, we find that debt relief is associated with a significant improvement in choice consistency. The proportion of participants with consistent responses in the money-earlier-or-later experiment rose significantly from 83.16% to 89.80% after debt relief (Two-tailed, two-sample test of proportions: z =-1.92, $p < 0.0548$). Because these improvements in consistency might be partially confounded with training effects due to pre-debt relief exposure to the experiment, we conduct a regression analysis of the determinants of choice consistency.

This Table reports results of fixed effects models where the dependent variable is whether the participant reports consistent choices in the money-earlier-or-later experiment (1: Yes, 0: No). We examine whether debt relief induced changes in psychological functioning (Columns 1-3), and changes in debt structure (Column 4), are associated with these post-relief improvements in choice consistency. The results in (Columns 1 and 3) suggest that participants with greater cognitive functioning, post-relief, are significantly more likely to exhibit improvements in consistency in the money-earlier-or-later experiment. There is no effect of changes in anxiety (Column 2), nor of changes in debt structure (Column 4), on choice consistency. These effects are significant when compared to the literature. (*27*) finds that a one-quartile increase in cognitive ability is associated with an increase of 5 percentage points in the proportion reporting consistent choices in a time discounting experiment. Our results suggest that the mean Flanker Combined Score improvement among participants of 1.25 points will increase the probability of consistent choice by about 3 percentage points – equivalent to improving by half a quartile in cognitive ability.

Table S13. Comparing the effects of financial resource shocks across studies

Panel A: Effects of Financial Resource Shocks on Cognitive Functioning and Risk Aversion

	Cognitive Functioning Response Time			CRRA Parameter	
	Our Study	Mani et al.	Carvalho et al.	Our study	Carvalho et al.
Before Debt Relief	0.264***			0.934***	
	(0.043)			(0.272)	
Before Harvest		0.19***			
		(0.036)			
Before Payday			0.02		-0.10
			(0.029)		(0.152)
Constant	3.379***	7.49***	8.06***	0.940***	1.66***
	(0.022)	(0.011)	(0.031)	(0.184)	(0.110)
Observations	392	902	20,206	392	1064
Number of Participants	196	451	1056	196	532

Panel B: Effect of Debt Relief on Cognitive Functioning using the Carvalho et al (6) Econometric Model

	Cognitive Functioning Flanker Task	
	ln (Response Time)	Correct (1: Yes, 0: No)
Before Debt Relief	0.283***	-0.133***
	(0.039)	(0.015)
Constant	1.137***	0.959***
	(0.043)	(0.015)
Observations	7,840	7,840
Number of Participants	196	196

Panel A, Cognitive Functioning Response Time: "Our study" dependent variable is the log of Flanker uncensored total response time in seconds; fixed effects model with no control variables. "Mani et al." dependent variable is the log of the numerical Stroop task total response time in seconds; uses re-analysis in (*6*) with a fixed effects model with controls for calendar month. "Carvalho et al." dependent variable is log of individual Flanker trial response time in seconds from Study 1 in (*6*); OLS model with controls for trial order. Panel A, CRRA Parameter: "Our study" and "Carvalho et al." dependent variable is the CRRA parameter interval; interval regression model with no control variables. Panel B, Cognitive Functioning Flanker Task: The regression models follow (6), where each participant has 20 Flanker Task trial observations. Each column reports results from an OLS regression that includes a Flanker Task trial-specific dummy, and clusters standard errors at the participant level. *p<0.10, **p<0.05, ***p<0.01

In common with our study, (*6*) and (*28*) measure the effect of financial resource shocks on cognitive functioning and risk aversion. Our study and (*6*) use the Eriksen Flanker Task to measure cognitive function, while (*28*) uses the Stroop Task. However, both the Flanker and the Stroop Task measure the inhibition control aspect of executive cognitive functioning (*29-30*). Our study and (*6*) use the same risk choice lottery experiment to measure risk aversion; the nominal values of the lottery are identical.

Panel A reports the effect of financial resource shocks across studies. The financial shocks are debt relief for chronically indebted low-income Singaporeans (our study), the annual harvest for Indian sugarcane farmers (*28*), and payday for the U.S. urban poor (*6*). The first two columns in Panel A show that our debt relief induced improvement in cognitive functioning is larger in magnitude than (*28*), and similar in precision. In contrast, even though we use the same Flanker task measure as (*6*), the improvements in cognitive functioning in our participants are more than ten times larger. The last two columns of Panel A compare changes in CRRA in our study and (*6*). Debt relief significantly reduces CRRA, making participants less risk-averse. In contrast, (*6*) find no payday effect on risk aversion.

Panel B addresses concerns that our results on cognitive functioning may differ from (*6*) simply because of econometric differences. Our measure of cognitive functioning in the main analysis and in Panel A is at the participant level; the main analysis uses the log of median response time, while Panel A uses the log of total response time for comparability with (*28*). However, (6) measures cognitive functioning at the individual Flanker Task trial level, generating 20 trial observations per subject in their sample (The actual average number of trials in their sample is slightly less than 20, as some of their subjects failed to complete all 20 trials). To demonstrate that our results are robust to these measurement differences, we re-estimate the effect of debt relief on cognitive functioning measured at the individual Flanker trial level (6).

We find broad support for the results reported in our main analysis: Average response time is significantly longer, and the probability of answering a trial correctly significantly lower, before debt relief compared to after debt relief. Our results appear robust to alternative methods of aggregating and estimating the cognitive functioning measurement.

Table S14. Questions to measure Generalized Anxiety Disorder (GAD)

Qn	Question Text	Response
1079	During the past 6 months, were you worried more than half the time?	Yes/No
1080	Did you find it difficult to control your worry?	Yes/No
Proceed to Q1081 to Q1086 if Q1079=Yes and Q1080=Yes		
Qn	**When you are worried or anxious, were you also…**	**Response**
1081	Restless or on edge?	Yes/No
1082	Particularly irritable?	Yes/No
1083	Aware of your muscles tensing?	Yes/No
1084	Easily tired?	Yes/No
1085	Having trouble falling asleep, or having restless unsatisfying sleep, or have trouble staying asleep?	Yes/No
1086	Having difficulty concentrating or did your mind go blank?	Yes/No

Question numbers are from the original survey.

Table S15. CRRA parameters in Eckel-Grossman risk aversion task

Option	Heads Payoff	Tails Payoff	CRRA Parameter Interval
1	\$28	\$28	$3.46 < r < +\infty$
2	\$36	\$24	$1.16 < r < 3.46$
3	\$44	\$20	$0.71 < r < 1.16$
4	\$52	\$16	$0.50 < r < 0.71$
5	\$60	\$12	$0 < r < 0.50$
6	\$70	\$2	$-\infty < r < 0$

Table S16. CRRA analysis with alternate lower CRRA bound

	Baseline CRRA Boundary r = -1	Lower CRRA Boundary r = -5
Debt Relief Amount	-0.110	-0.177*
	(0.072)	(0.094)
Debt Accounts Paid Off	-0.220	-0.237
	(0.140)	(0.167)
Observations	196	196
Number of Participants	196	196

See main paper Table 2; regression models are identical.

The lower bound assumption for the CRRA parameter affects the interval regression model estimates of the effect of debt relief on risk aversion, because a more negative lower bound means a change in risk task choices implies a much larger change in the participant's r. We examine how the estimates change when we use a significantly lower value of r = -5 as the CRRA lower bound, in place of the baseline estimates' value of r = -1. r = -5 implies a decision maker will be indifferent between accepting a choice that pays SGD28 with certainty, and a lottery that pays SGD32 or SGD0 with equal probability. The results show that making the CRRA lower bound more negative increases the magnitude and significance of the coefficient on debt relief amount. However, we do not believe these results have practical significance, because the implied risk choices with CRRA values lower than r = -1 are far riskier with those documented in the literature.

Table S17. Time discounting choice task

Decision Number	Option 1	Option 2
First Set (payment today vs. one month)		
Decision 1	$49 today	$50 in a month
Decision 2	$48 today	$50 in a month
Decision 3	$46 today	$50 in a month
Decision 4	$44 today	$50 in a month
Decision 5	$41 today	$50 in a month
Decision 6	$38 today	$50 in a month
Decision 7	$34 today	$50 in a month
Decision 8	$30 today	$50 in a month
Second Set (payment in 6 months vs. 7 months)		
Decision 9	$49 in 6 months	$50 in 7 months
Decision 10	$48 in 6 months	$50 in 7 months
Decision 11	$46 in 6 months	$50 in 7 months
Decision 12	$44 in 6 months	$50 in 7 months
Decision 13	$41 in 6 months	$50 in 7 months
Decision 14	$38 in 6 months	$50 in 7 months
Decision 15	$34 in 6 months	$50 in 7 months
Decision 16	$30 in 6 months	$50 in 7 months

Figure S1: Debt relief received by debt accounts paid off

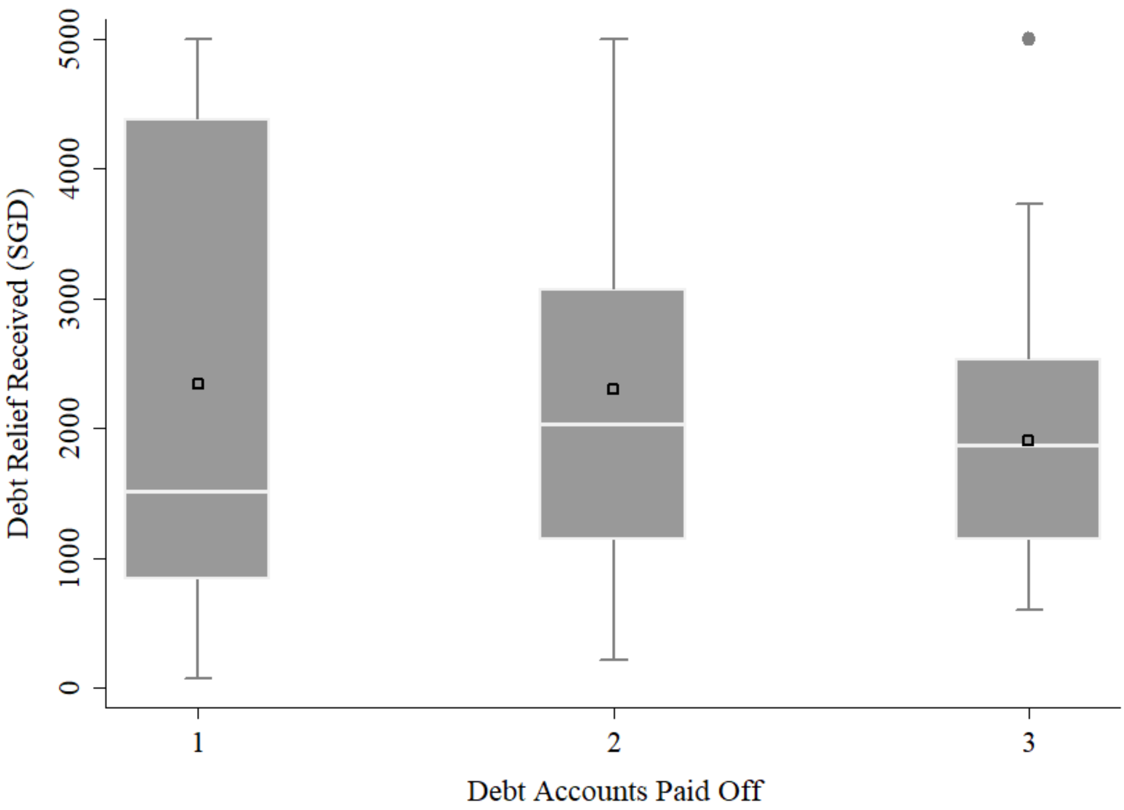

Based on final study sample. Mean debt relief received is represented by black square markers; the interquartile range of debt relief received is represented by the shaded box. Participants with zero debt accounts paid off received exactly SGD5000 in debt relief, and are not shown. A small number of participants (n=3) had four debt accounts paid off, and are not shown.

This Figure shows the distribution of debt relief received by the number of debt accounts (not debt categories) paid off, in the final study sample. As noted in the main paper, there is little collinearity between debt accounts paid off, and debt relief received, allowing the effects of each variable to be estimated in the same regression model. In fact, there is a slight negative correlation (-0.36) between debt accounts paid off and debt relief received. As there is also a slight positive correlation (0.3) between the number of initial debt accounts owed by participants, and their initial aggregate debt amount, this suggests social workers tended to prioritize larger debt accounts for repayment when endorsing clients with higher initial debts. As the larger accounts are generally housing related, this may reflect concerns from social workers that their clients might be made homeless if the housing debt accounts are not settled, or beliefs that smaller debts (but not larger debts) are potentially repayable through the client’s own efforts.

Figure S2. Training effects in cognitive functioning pre and post debt relief

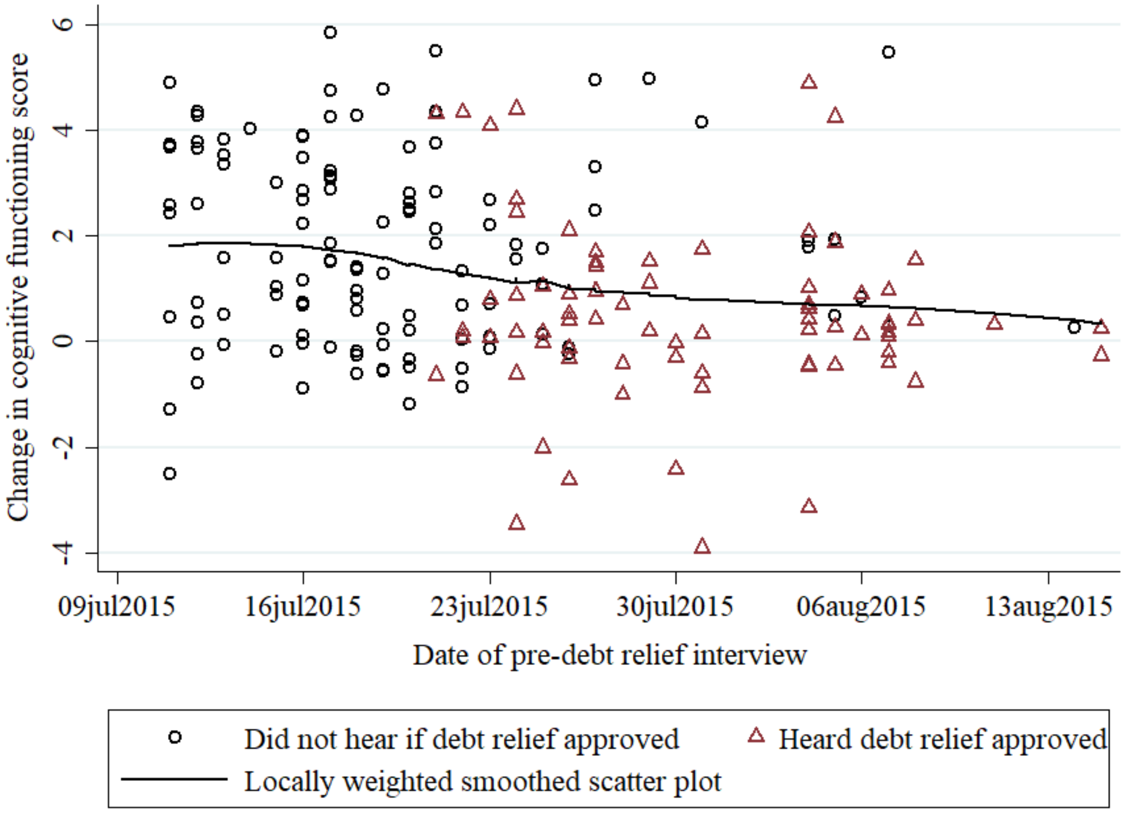

Training effects may explain why participants improve cognitive functioning, as they may have learned to respond to the Flanker task more quickly and accurately. In the main paper, we note that training effects alone would not explain why the structure of debt relief granted is strongly linked to the magnitude of improvements in cognitive functioning. Here, we present two analyses that further examine whether training effects are a major explanation of the improvement in cognitive functioning among beneficiaries of debt relief.

We conducted a simulation analysis of whether training effects are likely to be a strong factor behind our results. The NIH Toolbox Flanker Task validation study found a statistically significant and positive training effect over a two-week test-retest interval in the general U.S. adult population (*5*). While the population and test-retest interval are not ideally matched to our study, we are not aware of more appropriate published evidence from a large sample, nor were we able to conduct a test-retest exercise during our original study on a comparable group in Singapore. We therefore use (*5*) as a basis for simulating potential training effects, although we caution that (i) the two week test-retest interval is significantly shorter than the three month interval in our study, and (ii) the training effects are derived from a self-selected research volunteer population with higher mean cognitive performance than our community participants. These factors suggest the training effects in (*5*) may represent an upper bound.

In the simulation, we created 100,000 synthetic datasets by applying the training effect found in (*5*) to our data. Specifically, we first rescaled our participants' Flanker composite scores to a scaled score according to the conversion table for the NIH Toolbox Flanker measure in (*31*). The scaled score represents the performance of the participant compared to the average individual in the U.S. population, normed to a mean of 10 and standard deviation of 3. Our participants' pre-relief mean scaled score was 5.22 indicating performance significantly below the U.S. population average. The simulated training effect is added to each participant's scaled score, by drawing from a normal distribution: (mean = 0.79, SD = 1.73; as reported in (*5*)). We repeat this for each of the 196 participants and create 100,000 synthetic datasets. Each synthetic dataset thus represents a random draw from a potential counterfactual where all changes in performance in the Flanker Task are from training effects alone.

The simulated training effect produces a mean scaled score of 6.01, with the top percentile at 6.30. Both the mean and the top percentile scores are significantly below the participants' post-debt relief mean scaled score of 7.30. In 100,000 training effect counterfactuals, we were unable to replicate our study's mean post-debt relief improvements in cognitive functioning. While we acknowledge that training effects may contribute to enhanced Flanker Task performance post-debt relief, the best available evidence from the literature suggests it is very unlikely that training effects explain even half of the average improvement we find.

We also examined how improvements in cognitive functioning are affected by debt relief program anticipation effects. Anticipation effects may improve cognitive functioning and other outcome measures even before debt relief is actually paid, if participants have credible reasons to believe that their debts will be paid off soon. To minimize anticipation effects, we intended to complete first interview fieldwork prior to notification of debt relief approval. However, our plans were affected by an unexpected change in the debt relief program schedule. Program administrators informed us at the start of the study that they planned to announce the debt relief application results just before Singapore's National Day (9 August 2015). Two weeks into the survey period, program administrators informed us that results would be announced starting in late July. In response, we immediately introduced a new question to the survey on 21 July 2015 which asked survey participants if they had heard the results of their debt relief applications. Participants who took the survey later are more likely to have heard of the success of their debt relief application. Overall, 38% of the sample found out their applications had been approved by the time of their first interview – although debt relief payments were only issued from late August 2015 onwards. In addition, creditors operate on regular monthly billing cycles, so participants would only have received proof of debt relief when receiving their August or September billing statements. We include these participants in our main sample because anticipation effects, to the extent that they immediately relieve scarcity concerns, make it harder for us to detect any effects of debt relief. Indeed, the income shocks studied in the literature are completely anticipated, such as annual harvest income (*28*) or payroll (*6*).

This Figure plots the change in cognitive functioning of survey participants by the date of their pre-debt relief interview. Participants who were interviewed at a later date – and who had already heard of the debt relief application results – tended to have smaller improvements in cognitive functioning. This is consistent with the presence of anticipation effects of debt relief. While this does not rule out the presence of training effects, it does suggest that the changes in cognitive functioning we find are linked to debt relief itself.

Figure S3. Liquidity constraints: Informal credit access by debt accounts paid off

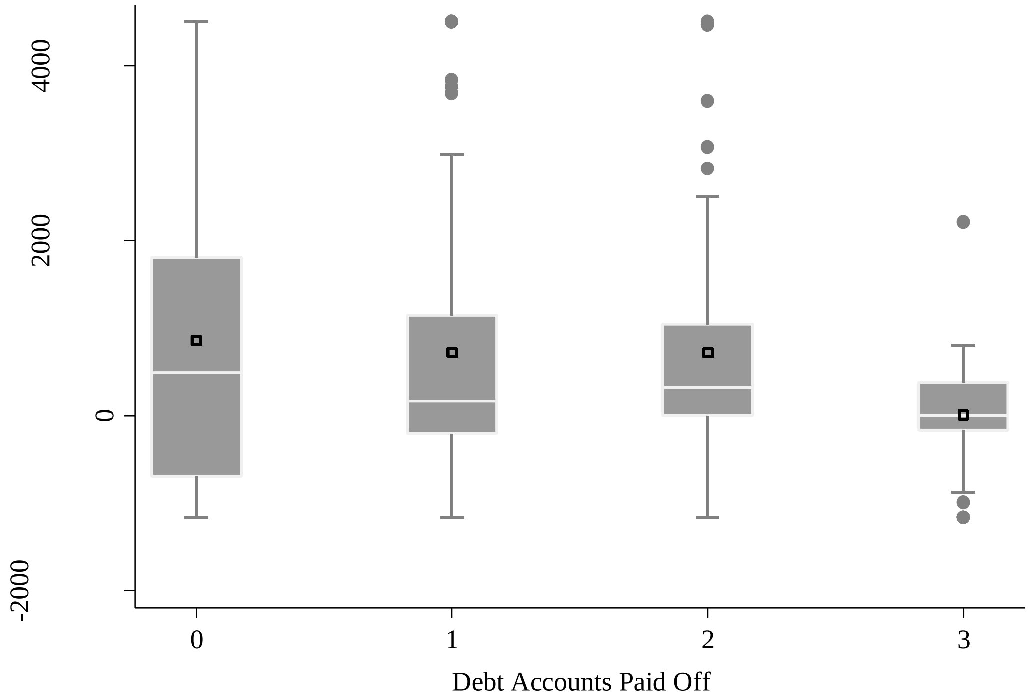

Informal Credit Access is censored at the 5th and 95th percentile. Mean consumption is represented by black square markers. A small number of participants (n=3) had four debt accounts paid off, and are not shown.

As discussed in the main paper, the debt relief program does not alter conventional financial liquidity because (i) low-income households are excluded from financial institution credit by Singapore regulators and (ii) the debt accounts paid off are generally not captured in Singapore consumer credit reports. However, beneficiaries might access informal credit by deferring monthly payments for creditors whose accounts were cleared by debt relief. For example, a beneficiary who had their utility debt paid off by debt relief could use the utility account for "liquidity" by not making full payment for their next monthly bill.

We estimate the potential value of informal credit access by measuring the difference between the size of debt relief granted, and the change in actual debt three months post-relief. That is, Informal Credit Access = (Post-Debt Relief Debt Size – Pre-Debt Relief Debt Size) – (Size of Debt Relief Grant). Therefore, a beneficiary who received SGD5000 in relief, but only has a change in actual debt of SGD4000, has informal credit access worth SGD1000.

This Figure reports informal credit access by the number of debt accounts paid off during debt relief. On average, informal credit access is worth SGD556. There is no strong correlation between debt accounts cleared and informal credit access, suggesting that greater intensity of debt relief, as measured by clearing more debt accounts, is not correlated with greater informal credit access. Therefore, informal credit access does not explain why our main analysis shows a greater number of debt accounts cleared drives greater improvements in cognitive functioning and changes in economic decision-making.

Figure S4. Screenshots of Eriksen Flanker task

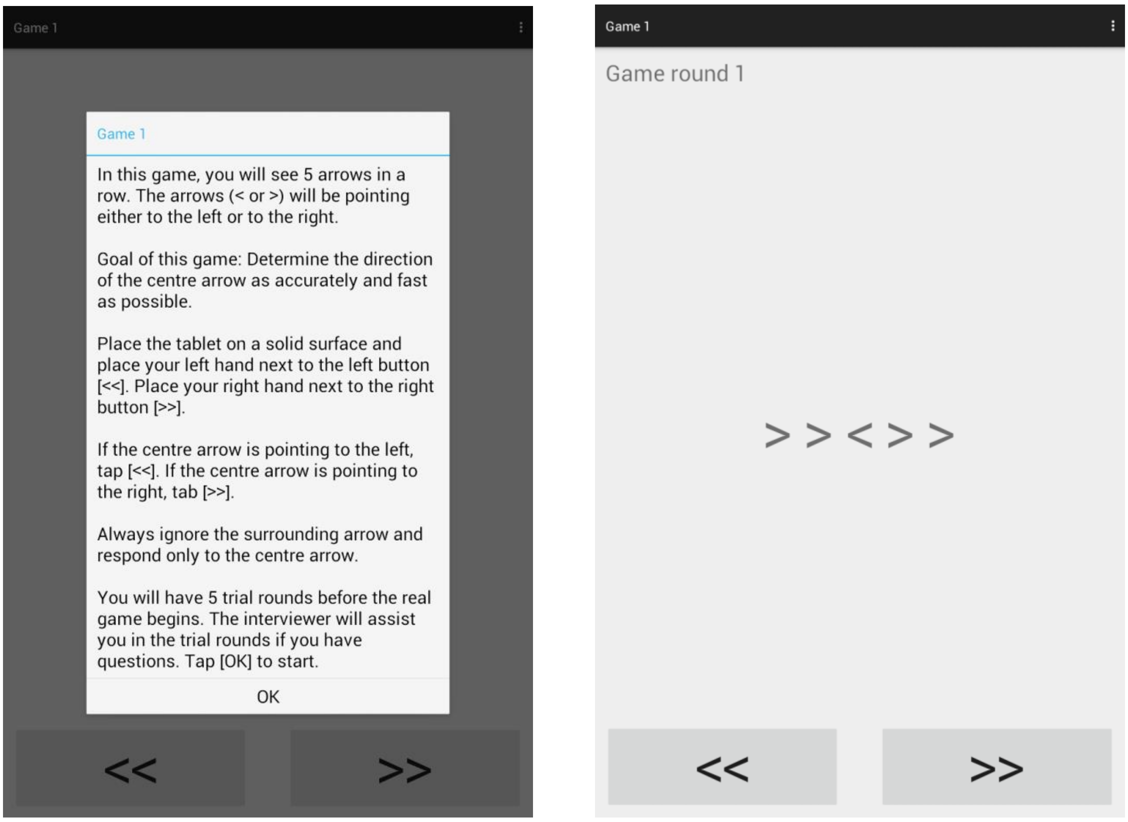

Figure S5. Screenshots of Eckel-Grossman risk aversion task

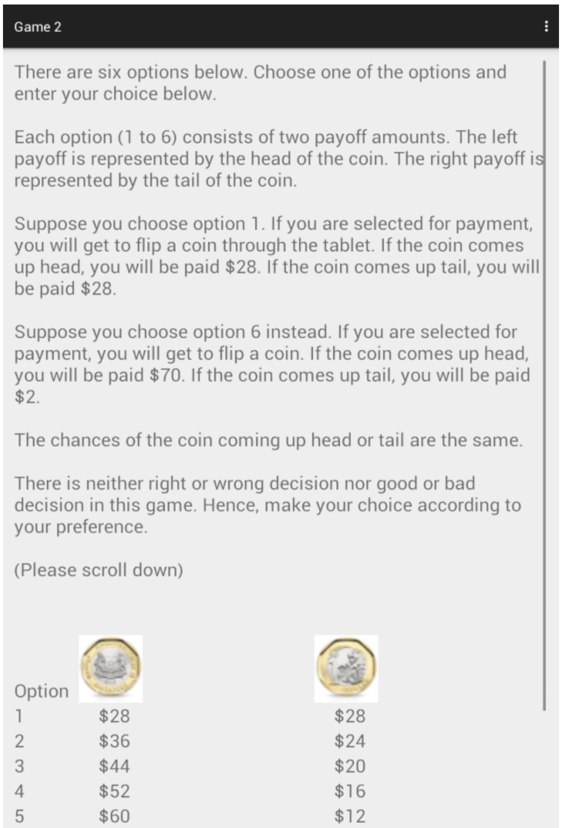

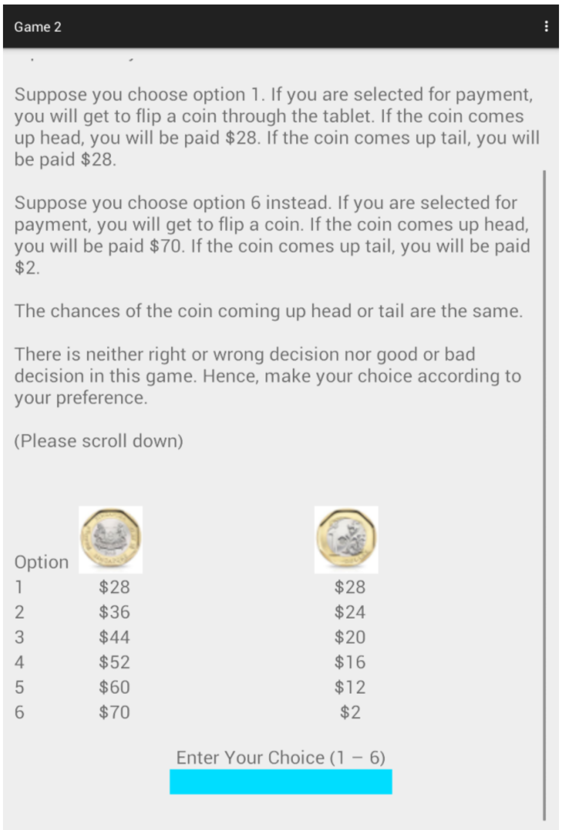

Figure S6. Screenshots of time discounting task

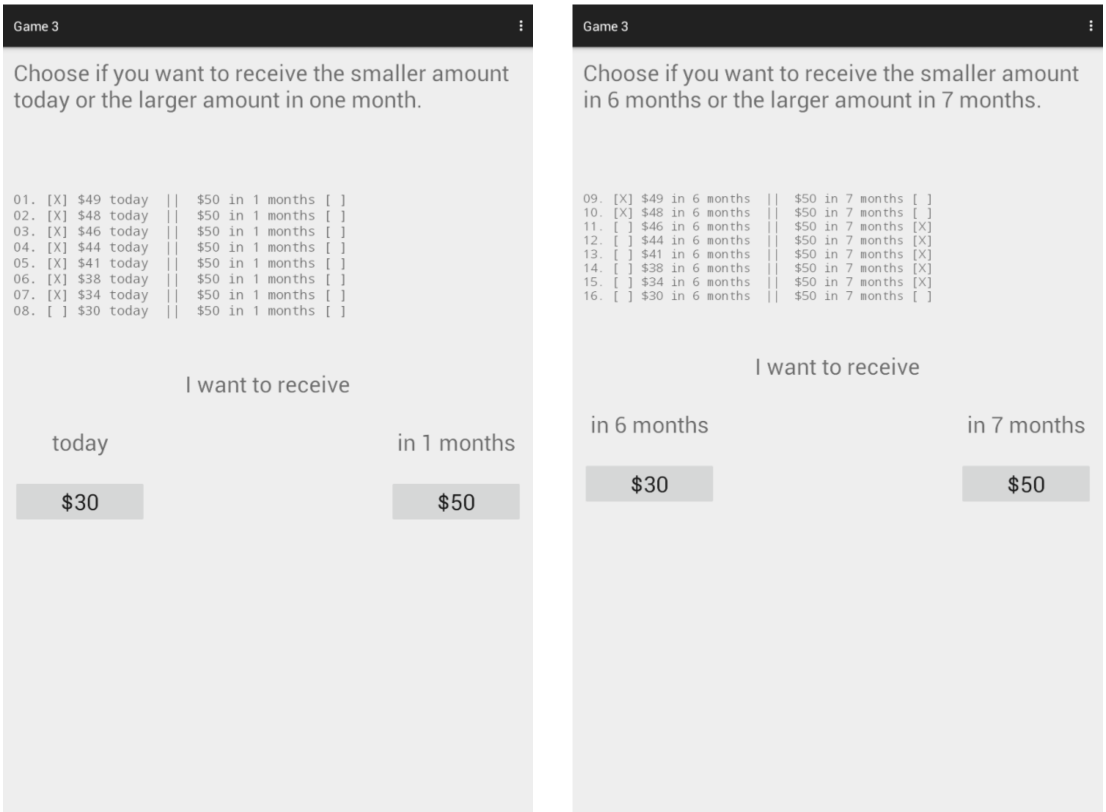

SI Appendix References:

1. L.C. Johnson, S.J. Yanca. Social Work Practice: A Generalist Approach (8th Edition) NY: Pearson (2004).
2. Ministry of Social and Family Development. (2015). Volume 1: Working with Vulnerable Families: Practitioner's Resource Guide. Accessed on 1 Nov 2018 at: https://www.msf.gov.sg/publications/Pages/Strengthening-Families-Together-SFT-Pilot-Practitioners-Resource-Guides.aspx.
3. I.Y.H. Ng, Poverty Attitudes of Singaporeans: A Question of Class, Politics, and Action? *Social Indicators Research* **121**, 371-385 (2015).
4. J. Slotkin, C. Nowinski, R. Hays, J. Beaumont, J. Griffith, S. Magasi, J. Salsman, R. Gershon, NIH Toolbox® Scoring and Interpretation Guide. Chicago, IL: Northwestern University and National Institutes of Health (2012).
5. P. D. Zelazo, J. E. Anderson, J. Richler, K. Wallner-Allen, J. L. Beaumont, S. Weintraub, II. NIH Toolbox Cognition Battery (CB): Measuring executive function and attention. *Monographs of the Society for Research in Child Development* **78**, 16-33. (2013).
6. L. S. Carvalho, S. Meier, S. W. Wang, Poverty and economic decision-making: Evidence from changes in financial resources at payday. *American Economic Review* **106**, 260-284 (2016).
7. M. R. Liebowitz, Generalized anxiety disorder (includes overanxious disorder of childhood). in DSM-IV sourcebook volume, T. A. Widiger, Ed.. (Washington, DC: American Psychiatric Association,1996).
8. J. C. Baer, M. Kim, B. Wilkenfeld, Is it generalized anxiety disorder or poverty? An examination of poor mothers and their children. *Child and Adolescent Social Work Journal* **29**, 345-355 (2012).
9. C. Eckel, P. J. Grossman, Sex differences and statistical stereotyping in attitudes toward financial risk. *Evolution and Human Behavior* **23**, 281-295 (2002).
10. C. Dave, C. C. Eckel, C. A. Johnson, C. Rojas, Eliciting risk preferences: When is simple better? *Journal of Risk and Uncertainty* **41**, 219-243 (2010).
11. R. Mehra, E. C. Prescott, The Equity Premium: A Puzzle. *Journal of Monetary Economics* **15**, 145-161 (1985).
12. C. A. Holt, S. K. Laury, Risk Aversion and Incentive Effects. *American Economic Review* **92**, 1644-1655 (2002).
13. G. W. Harrison, M. I. Lau, M. B. Williams, Estimating individual discount rates in Denmark: A field experiment. *American Economic Review* **92**, 1606-1617 (2002).
14. S. Meier, C. Sprenger. Present-biased preferences and credit card borrowing. *American Economic Journal: Applied Economics*, *2*, 193-210 (2010).
15. J.D. Cohen, K.M. Ericson, D. Laibson, J.M. White. Measuring Time Preferences. National Bureau of Economic Research Working Paper 22455 (2016).
16. J. Andreoni, C. Sprenger, Risk preferences are not time preferences. *American Economic Review*, **102**, 3357-3376 (2012).
17. D. Laibson, Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, **112**, 443-478. (1997).
18. M.E. Schaffer, xtivreg2: Stata module to perform extended IV/2SLS, GMM and AC/HAC, LIML and k-class regression for panel data models. (2010) Accessed at: http://ideas.repec.org/c/boc/bocode/s456501.html
19. D. Roodman, Estimating fully observed recursive mixed-process models with cmp. Stata Journal 11(2): 159-206 (2011)

20. J. D. Angrist, G. W. Imbens, Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity. *Journal of the American Statistical Association* **90**, 431-442 (1995).
21. Housing Development Board Annual Report 2015/2016, Housing and Development Board, Singapore.
22. J. M. Hinson, T. L. Jameson, P. Whitney, Impulsive decision making and working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **29**, 298-306 (2003).
23. P. Whitney, C. A. Rinehart, J. M. Hinson, Framing effects under cognitive load: The role of working memory in risky decisions. *Psychonomic Bulletin & Review* **15**, 1179-1184 (2008).
24. T. Dohmen, A. Falk, D. Huffman, U. Sunde, Are risk aversion and impatience related to cognitive ability? *American Economic Review* **100**, 1238–1260 (2010).
25. D. J. Benjamin, S. A. Brown, J. M. Shapiro, Who is 'behavioral'? Cognitive ability and anomalous preferences. *Journal of the European Economic Association* **11**, 1231-1255 (2013).
26. J. Haushofer, E. Fehr, On the psychology of poverty. *Science* **344**, 862-867 (2014).
27. S.V. Burks, J.P. Carpenter, L. Goette, A. Rustichini. Cognitive skills affect economic preferences, strategic behavior, and job attachment. *Proceedings of the National Academy of Sciences*, 106, 7745-7750 (2009).
28. A. Mani, S. Mullainathan, E. Shafir, J. Zhao, Poverty impedes cognitive function. *Science* **341**, 976-980 (2013).
29. A. Diamond, Executive Functions. *Annual Review of Psychology* **64**, 135-168 (2013).
30. E. B. Dean, F. Schilbach, H. Schofield, "Poverty and cognitive function" in *The Economics of Asset Accumulation and Poverty Traps* (University of Chicago Press 2017).
31. K. B. Casaletto, A. Umlauf, J. Beaumont, R. Gershon, J. Slotkin, N. Akshoomoff, R. K. Heaton, Demographically corrected normative standards for the English version of the NIH Toolbox Cognition Battery. *Journal of the International Neuropsychological Society*, ***21***, 378-391 (2015).