MAX-PLANCK-INSTITUT FÜR SOZIALRECHT UND SOZIALPOLITIK
MAX PLANCK INSTITUTE FOR SOCIAL LAW AND SOCIAL POLICY

mea Munich Center for the Economics of Aging

Heterogeneous Effects of Poverty on Cognition

Helmut Farbmacher, Heinrich Kögel, Martin Spindler

06-2019

MEA DISCUSSION PAPERS

Heterogeneous Effects of Poverty on Cognition

Helmut Farbmacher, Heinrich Kögel, and Martin Spindler

Abstract:

We examine heterogeneity in the effect of poor financial circumstances on cognition. Our analysis uses data from an experiment, which randomly assigned low-income individuals to perform a cognitive test before or after payday. On average, and based on traditional subgroup analysis, the experiment did not suggest that the poorer financial circumstances before payday impeded cognitive function. Using the causal forest method, however, our heterogeneity analysis suggests that there are indeed detrimental effects among young and elderly individuals with very low incomes. We can confirm this finding in an independent experiment, using only traditional subgroup analysis.

Zusammenfassung:

Wir untersuchen, ob schlechte finanzielle Umstände heterogene Effekte auf kognitive Funktion haben. Hierfür haben wir Daten aus einem Experiment ausgewertet, in welchem einkommensschwache Individuen einen kognitiven Test vor oder nach Erhalt ihrer Löhne absolvierten. Das Experiment lieferte im Durchschnitt und auf Basis einer gewöhnlichen Subgruppenanalyse keine Anzeichen dafür, dass schlechtere finanzielle Umstände vor der Lohnzahlung kognitive Funktion beeinträchtigten. In unserer Heterogenitätsanalyse, die wir mit Hilfe der Causal Forest Methode durchführen, finden wir allerdings nachteilige Effekte unter jüngeren und älteren Individuen mit niedrigem Einkommen. Wir können diese Ergebnisse durch eine gewöhnliche Subgruppenanalyse in einem unabhängigen Experiment bestätigen.

Keywords:

Poverty, cognition, heterogeneous effects, causal forest

JEL Classification:

D14, D91, I32

Heterogeneous Effects of Poverty on Cognition*

Helmut Farbmacher[1,3], Heinrich Kögel[1,3], and Martin Spindler[2,3]

[1] Department of Economics, University of Munich (LMU), Munich, Germany
[2]Hamburg Business School, University of Hamburg, Hamburg, Germany
[3]Munich Center for the Economics of Aging (MEA), Max Planck Society, Munich, Germany

December 9, 2019

Abstract

We examine heterogeneity in the effect of poor financial circumstances on cognition. Our analysis uses data from an experiment, which randomly assigned low-income individuals to perform a cognitive test before or after payday. On average, and based on traditional subgroup analysis, the experiment did not suggest that the poorer financial circumstances before payday impeded cognitive function. Using the causal forest method, however, our heterogeneity analysis suggests that there are indeed detrimental effects among young and elderly individuals with very low incomes. We can confirm this finding in an independent experiment, using only traditional subgroup analysis.

Keywords: poverty, cognition, heterogeneous effects, causal forest

JEL-Codes: D14, D91, I32

*We thank Axel Börsch-Supan, Martin Huber, Mike Hurd, Jann Spiess, Stefan Wager, Frank Windmeijer, Joachim Winter, and participants of the Machine Learning in Economics and Econometrics Workshop in Munich 2018, the Annual Conference of the German Health Economics Association in Augsburg 2019, the 6th Annual Conference of the International Association for Applied Econometrics in Nicosia 2019, the 31st Conference of the Austro-Swiss Region of the International Biometric Society in Lausanne 2019, as well as seminars at Hamburg University, Harvard University, MEA and TU Braunschweig for helpful comments. The authors declare that they have no relevant or material financial interests related to the research described in this paper.

1 Introduction

Many studies have documented associations between poverty and less beneficial behavior. For example, the poor are less likely than those with higher incomes to make use of preventive health services, and more likely to smoke cigarettes, play the lottery, and borrow more often at high cost.[1] Despite long-standing debates in economics and other disciplines, the reasons for such behavior remain unclear and the topic itself controversial. One recent hypothesis has focused on the financial circumstances of the poor and the potentially detrimental impact of these on cognition: In a sample of farmers from India, Mani et al. (2013) found that participants showed reduced cognitive performance before harvest, when poor, compared to after harvest, when rich. The authors suggested that a preoccupation with monetary concerns may leave the farmers before harvest with fewer mental resources available for other processes.[2]

In the only other study to have investigated this hypothesis empirically to date, Carvalho et al. (2016) assigned a sample of low-income US individuals randomly to perform a number of cognitive tests before or after payday. The individuals surveyed before payday faced poorer financial circumstances than those surveyed after payday. However, the authors found no before-after differences in cognitive function in the full sample or selected subgroups. These mixed empirical findings, and the dearth of studies on this hypothesis in general, highlight the need to identify, at a more detailed level, the groups of individuals in which poor financial circumstances might have detrimental effects on cognitive function.

To contribute to this area of study, we therefore analyze heterogeneity in the effect of financial circumstances on cognition, focusing on identifying individuals in whom poorer financial circumstances have negative effects. To do so, we use data from the experiment conducted by Caravalho et al. (2016). For our heterogeneity analysis, we use the causal forest method by Athey et al. (2019), which was developed specifically to explore heterogeneous treatment effects in experiments. The method can be described as an adaptive nearest-neighbors approach that exploits ideas from the random forest machine learning literature to determine the relevant neighborhoods for estimating conditional average treatment effects at given points in the covariate space. Compared with traditional ordinary least squares (OLS) subgroup analyses, the causal forest method allows non-linear treatment effects to be esti-

[1] Use preventive health services (Ross et al. 2007), smoke cigarettes (Dube et al. 2009), play the lottery (Clotfelter et al. 1999), borrow at high cost (Bourke et al. 2012).

[2] See Bertrand et al. (2004; 2006) for a discussion of alternative views on the behavior of the poor.

mated in a fully flexible way and circumvents the need to specify an interacted model, which may not always be straightforward (especially when the number of covariates is large). We examine effect heterogeneity using a rich set of 37 policy-relevant, pre-treatment covariates, including age, income, employment status, and measures of financial strain in the past. Our causal forest analysis proceeds in the following steps: First, we investigate which covariates are particularly relevant for heterogeneity in the treatment effect. Next, we examine how the effect varies across the most important variables. Subsequently, we study, in greater detail, the effect heterogeneity in regions of the covariate space where the previous step indicates particularly detrimental effects.

The results of our analysis suggest that there is strong effect heterogeneity in the two covariates age and income. For old and young individuals who received a very low income around the time of the experiment, we find that the poorer financial circumstances before payday had detrimental cognitive effects. We verify this finding using a second, independent, experiment conducted by Carvalho et al. (2016). Our results provide further evidence that there may be a causal effect of poverty on cognition. They also demonstrate the benefit of using the causal forest method to identify treatment effect heterogeneity that may have been overlooked in traditional subgroup analyses.

The remainder of this paper is structured as follows. Section 2 describes the experiment and our analysis sample. Section 3 explains the causal forest method. Section 4 presents average effect estimates for the full sample, the results of our heterogeneity analysis, and investigates the findings of our heterogeneity analysis in an independent experiment. Section 5 concludes.

2 Experiment and Data

2.1 Experiment

Carvalho et al. (2016) conducted their experiment twice, once among members of the RAND American Life Panel and then again among members of the GfK KnowledgePanel. Both are ongoing online panels with individuals aged 18 and over living in the United States. The authors restricted the sample for each experiment to individuals with an annual household income of $40,000 or less. For our analysis, we use the data from the GfK KnowledgePanel because it had the larger sample size, and because its share of compliers, i.e. the proportion of individuals who actually completed the survey before payday out of all individuals assigned to the before-payday group, was much higher. The following descriptions therefore pertain to the GfK KnowledgePanel.

The experiment consisted of a baseline survey and a follow-up survey, the former of which was used to determine individuals' paydays and the latter of which was used to administer the cognitive test. Individuals were randomly assigned to receive the survey with the cognitive test before or after payday.

In the baseline survey, individuals were asked to state all of the dates and amounts of payments that they (and their spouse) expected to receive during a reference period from 21 November to 20 December 2014. All individuals who did not give full information about the number and dates of expected payments, or who reported expected payments for more than two different dates, were dropped from the sample.[3] Using this payment information, Carvalho et al. (2016) defined each individual's payday as follows: For individuals whose largest payment arrived at least two weeks after the previous payment, the date of the largest payment was set as the payday. For all other individuals, the payday was determined to be the payment date after the longest period without payment. If an individual's payments were fewer than two weeks apart, he or she was also excluded from the experiment.

The follow-up survey opened one week before payday for individuals assigned to the before-payday group and one day after payday for individuals assigned to the after-payday group. Carvalho et al. (2016) found that 98 percent of all individuals assigned to be surveyed before payday actually completed the survey before payday. Despite this high compliance rate, we follow Carvalho et al. (2016) in our analysis and estimate intention-to-treat effects, using the random assignment to the before-payday group as the regressor of interest.

The cognitive test in the follow-up survey was a version of the numerical Stroop task, which measures cognitive control. Participants are shown a number that consists of a repeated digit (e.g., 555). Subsequently, they must state, as quickly as possible, how many times the digit is repeated in the number rather than stating the digit itself – the correct answer in the example being three rather than five. The experiment by Carvalho et al. (2016) ran the Stroop task with 48 trials, and per trial each individual had, at most, five seconds to respond – otherwise the answer to the trial was coded as incorrect.

To confirm that the individuals actually experienced poorer financial circumstances be-

[3] The latter restriction was imposed to remove individuals for whom consumption smoothing may be easier.

fore payday than they did after payday, the follow-up survey also collected information on individuals' cash holdings, checking and savings accounts balances, and total expenditures over the past seven days. Based on these measures, Carvalho et al. (2016) showed that the experiment had indeed created substantial variation in financial circumstances.[4] Appendix Table 1 presents results from our estimations that are analogous to Carvalho et al.'s (2016) for financial circumstances. These estimations yield very similar variation in financial circumstances in our sample, which is slightly smaller than Carvalho et al.'s (2016) sample, as explained in the next section.

2.2 Sample and Descriptive Statistics

For our analysis sample, we select all of the 2,723 individuals who were in Carvalho et al.'s (2016) full KnowledgePanel sample and subsequently drop all observations that are missing information on any of our analysis variables.[5] This selection procedure yields a sample of 2,480 individuals.

Table 1 presents the definitions and descriptive statistics for the cognition outcomes and treatment indicator. Our main outcome of interest is the number of correct answers per second that individual i gave over the entire Stroop task. This outcome captures the essence of the Stroop task's goal, which is to give correct answers to all trials as quickly as possible. Moreover, to gain an understanding of where the effect on our main outcome comes from, we include the numerator and denominator of our main outcome as additional outcomes: the number of correct answers over all 48 trials and the total time it took individual i to complete the entire Stroop task.[6]

Table 1 shows that, on average, the individuals in our sample gave approximately 0.45 correct answers per second, provided about 43 correct answers in total (thereby responding correctly to most of the trials), and took approximately 100 seconds to finish the whole Stroop task. The mean for our regressor of interest, which is a dummy that is equal to 1 if an individual was randomly assigned to be surveyed before payday and 0 otherwise, is almost

[4] This finding is in line with previous research, which documented a sharp increase in caloric intake and expenditures at payday for certain groups of individuals (see, for example, Mastrobuoni and Weinberg 2009, Shapiro 2005, Stephens 2003, Stephens 2006).

[5] Additionally, we drop all individuals who were above the 0.99 quantile of the current income distribution in our full sample to remove potentially erroneous values. Given the definitions of our outcomes below, we also drop individuals who have missing information for any of the Stroop task's trials, i.e., who did not participate in all 48 trials of the task.

[6] Carvalho et al. (2016) conducted their Stroop task analysis at the individual x trial level, using the outcomes response time per trial and a dummy which is one if an individual answered a trial correctly.

exactly 50 percent. This is as expected considering the experiment's random assignment of individuals to the before-payday or after-payday group.

Table 2 reports descriptive statistics for the 37 covariates that we include in our heterogeneity analysis. All of these were collected before the follow-up survey, in which the Stroop task was administered.[7] These covariates give information on many policy-relevant characteristics, such as an individual's race, education, employment status, and financial strain in the past. In addition to the annual household income at the time of the baseline survey, we include a measure of the (household) income that an individual received around the time of the experiment. We call this measure current income and construct it as the sum of all payments that an individual (and his or her spouse) expected to receive during the experiment's reference period (21 November to 20 December 2014). Overall, Table 2 suggest that many individuals in the sample were of low socioeconomic status. For example, 41.4 percent of them had experienced financial hardship in the past 12 months, and almost half stated that they were living from paycheck to paycheck. Also, the annual household income dummies show that 41.1 percent of all individuals had an annual household income of less than $20,000, and an average current income of approximately $1738.

3 Methodology

The goal of our analysis is to study heterogeneity in the effect on cognition of poorer financial circumstances before payday. To do so, we estimate conditional average treatment effects using the causal forest method, which is based on the generalized random forest framework by Athey et al. (2019). The method is designed for studying treatment effect heterogeneity in experiments and can be described as an adaptive nearest-neighbors approach that uses a type of random forest technique to determine the weighting of observations in the estimation procedure.[8] This section describes the main idea of the causal forest. For technical details, see Athey et al. (2019).

To fix ideas, assume the following random effects model for individual i, $i = 1, \ldots, n$:

$$Y_i = \tau_i D_i + \epsilon_i, \tag{1}$$

[7]Appendix Table 2 shows that the experiment's randomization procedure was successful in balancing the analysis covariates between the individuals interviewed before and after payday.

[8]For an introduction to random forests, see, for example, Hastie et al. (2009).

where Y_i is one of our cognition outcomes, ϵ_i is i's outcome when assigned to be surveyed after payday, D_i is a dummy that equals one if individual i was assigned to be surveyed before payday, and τ_i corresponds to the effect of the financial circumstances before payday for individual i. Due to the random assignment of individuals to the before-payday or after-payday group, it further holds that D_i is independent of τ_i and ϵ_i.

Our quantity of interest is the conditional average treatment effect $\tau(x) = E(\tau_i|X_i = x)$, which in our case is the average effect of the financial circumstances before payday on cognition at a point x of the covariate vector X_i. For the estimation of $\tau(x)$, the causal forest method exploits the independence assumption of D_i and sets up two local moment equations. In the next step, the method obtains an estimate for $\tau(x)$ by fitting an empirical version of the local moment equations.[9] This procedure yields the causal forest estimator $\hat{\tau}$(x), which can be written as:

$$\hat{\tau}(x) = \sum_{\{i:D_i=1\}} \frac{\alpha_i(x)}{\sum_{\{i:D_i=1\}} \alpha_i(x)} Y_i - \sum_{\{i:D_i=0\}} \frac{\alpha_i(x)}{\sum_{\{i:D_i=0\}} \alpha_i(x)} Y_i, \tag{2}$$

where $\alpha_i(x)$ is a type of similarity weight, measuring individual i's relevance in the estimation of $\tau(x)$. Thus, the causal forest estimator estimates $\tau(x)$ by taking the difference in weighted average outcomes between the treated and untreated individuals.

To determine the weights $\alpha_i(x)$, the causal forest algorithm uses an approach that is based on the random forest method. The goal of Breiman's (2001) original random forest is to predict an outcome Y_i using covariates X_i by averaging over predictions from an ensemble of trees. Each tree is constructed by recursively splitting the covariate space into axis-aligned partitions, whereby at every step the split is chosen to maximize the tree's prediction accuracy. The prediction accuracy is typically evaluated using the mean squared error. After a stopping criterion has been reached, a single tree thus yields a partitioning of the covariate space into disjoint regions, or leaves, and its prediction for Y_i at point $X_i = x$ is calculated as the average Y_i over all observations that fall into the same leaf, based on their values in X_i, as the point x. For the construction of each tree, a different bootstrap sample of the data is used, and at every step only a random subset of all covariates is made available for splitting. Appendix B shows an example of a single tree.

Now, for obtaining the weights $\alpha_i(x)$, the causal forest also grows an ensemble of trees using recursive partitioning. However, rather than averaging over predictions from the trees,

[9]See Appendix A for details.

the causal forest counts how many times individual i is in the same leaf as point x across all constructed trees, and derives $\alpha_i(x)$ based on this number. Specifically, for a set of trees $b = 1, \ldots, B$, the weight $\alpha_i(x)$ for individual i is computed as follows:

$$\alpha_i(x) = \frac{1}{B} \sum_{b=1}^{B} \frac{1\{i \in I_b(x)\}}{n_b(x)}, \tag{3}$$

where $I_b(x)$ is the set of all indices for the individuals that are in the same leaf as point x in tree b, and $n_b(x)$ is the number of individuals that fall into the same leaf as x in tree b. Thus, the more often individual i is in the same leaf as point x, the more weight i receives in estimating $\tau(x)$.

Compared with the random forest algorithm described above, the causal forest also uses a different splitting criterion for constructing the trees. The causal forest criterion is based on treatment effect estimates within the covariate space partitions, and, at a high level, implies that the algorithm seeks to maximize the treatment effect heterogeneity across partitions at every tree-splitting step. Athey et al. (2019) show that maximizing this criterion is related to improving the tree's expected accuracy in predicting treatment effects (rather than the outcome Y_i) at every step of the splitting procedure.

The causal forest also only allows splitting at every step based on a random subset of the covariates. In addition, the algorithm grows its trees on random subsamples of the data and implements a subsample splitting technique Athey et al. (2019) call honesty.[10] The idea behind the honest approach is to split a given subsample randomly into two roughly equally sized parts. The tree structure is subsequently grown on one of the two subsample parts, and the resulting structure is used to determine which individuals in the other subsample part are in the relevant neighborhood for estimating $\tau(x)$. Intuitively, the approach implies that observation i's outcome Y_i is not able to influence the construction of its weight $\alpha_i(x)$. This guards against spuriously extreme Y_i values obtaining unduly large influence in the data-driven weight calculation and thereby confounding the estimate for $\tau(x)$.

Athey et al. (2019) show that the causal forest estimates are consistent and asymptotically normally distributed, and derive bootstrap standard errors that allow for constructing valid confidence intervals.

We conduct our analysis in R, using the package grf by Tibshirani et al. (2018). The package implements the causal forest estimator in the function causal_forest, and also in-

[10]See Athey and Imbens (2016) and Wager and Athey (2018) for discussions of honesty.

cludes the bootstrap standard errors.[11] We estimate three causal forests, i.e., one for each of our three outcomes. We grow each forest using 10,000 trees with at least two observations per leaf. Following the function's default values, we build each tree on a 50 percent subsample of our analysis sample, using the honest approach, and allow 27 of our 37 covariates as tree-splitting candidates at each step.[12]

4 Results

Section 4.1 describes the OLS average effect estimates for the full sample. Section 4.2 subsequently presents the results of our heterogeneity analysis, and Section 4.3 gives the estimates for our subgroup analysis based on the insights from the heterogeneity analysis, using our main analysis sample and an additional, independent, sample by Carvalho et al. (2016).

4.1 OLS Analysis

Table 3 displays the OLS estimates for the average effect of the financial circumstances before payday on the main outcome – i.e., the number of correct answers per second – and the two additional outcomes: number of correct answers and total response time. As can be seen in Column (1), the estimated effect on the number of correct answers per second is statistically insignificant at the 10 percent level, and the point estimate's magnitude of 0.007 appears small relative to the average number of correct answers per second for the after-payday group, which is 0.443. In addition, the sign of the effect point estimate goes in the direction opposite to that which one would expect if the poorer financial circumstances before payday were to impede cognitive function: on average, the individuals assigned to the before-payday group gave a greater number of correct answers per second than did the individuals assigned to the after-payday group. Similar to the results in Column (1), the estimations for the other two outcomes, shown in Columns (2) and (3), also yield effect estimates that are insignificant at the 10 percent level, small in magnitude, and whose signs go in the direction opposite to that which is expected.

[11] The function optimizes an approximation of the theoretically motivated tree-splitting criterion to increase computational efficiency. See Athey et al. (2019) for details.

[12] Because we use the honest approach in our estimation, effectively a 25 percent subsample is used for growing each tree. For the other parameters that need to be specified in the causal_forest function, we also use the function's default values, and we enable the local centering feature of the algorithm.

In short, the estimates in Table 3 do not suggest that, on average, the poorer financial circumstances before payday have a detrimental effect on cognition in the full sample. This finding is in line with Carvalho et al.'s (2016) results.[13]

4.2 Heterogeneity Analysis

Our heterogeneity analysis proceeds in three steps. First, we calculate a variable importance measure for our three causal forests to identify which of the 37 covariates may be especially important for heterogeneity in our effects of interest. Next, based on these insights, we investigate in heatmaps how the conditional average treatment effects vary over the two most important variables. Subsequently, we estimate effects for two 'typical' individuals in two regions in which the heatmaps suggest particularly detrimental effects, and study how the effect estimates change when we vary the values of the 35 remaining covariates.

To assess variable importance in our estimated causal forests, we use a measure implemented in the grf R package. For variable X_k, the variable importance measure essentially captures the relative frequency with which a forest split on X_k across all grown trees. The measure, therefore, gives an indication over which variables the conditional average treatment effect may vary the most. For X_k, the measure ranges from 0, if the forest never split on X_k, to 100, if the forest always split on X_k.[14] Panel A in Figure 1 shows the variable importance plot for the causal forest using the number of correct answers per second as the outcome. The panel yields that by far the two most important variables in the tree-splitting procedure are the covariates age and current income. Both have a variable importance value of approximately 25. All other covariates have a value of around five at most. Similarly, for the two causal forests using the outcomes number of correct answers and total response time, Panels B and C in Figure 1 also suggest that age and current income are by far the most important variables.

Next, to explore how the effects vary in age and current income, Figure 2 displays heatmaps, plotting effect estimates over an age–current income grid. The maximum value on the x-axis of \$1500 corresponds to the median current income in our sample. For estimating the effects, we set all other continuous and categorical covariates to their full sample median,

[13] Appendix Table 3 additionally shows our effect estimates for the subgroups analyzed by Carvalho et al. (2016). Also in line with their results, our estimations yield effect estimates that are insignificant at the 10 percent level and small in magnitude for all subgroups.

[14] See Appendix C for details.

and all dummy covariates according to the most frequently occurring characteristics in the full sample. For example, 76.1 percent of all individuals in the sample are white. Therefore, we set the dummy white equal to one, and all other race dummies to zero.[15] Red regions indicate effect estimates that are detrimental and blue regions indicate effect estimates that are not detrimental.

Panel A in Figure 2 displays the estimated effects for the number of correct answers per second. The panel shows that the causal forest estimates negative effects especially for individuals who have a current income below approximately \$750 and whose age is either up to approximately 30 years or between around 70 and 80 years. A current income of \$750 appears rather low, corresponding to the 0.16 quantile of our sample's current income distribution. For the younger individuals with a lower current income, the estimated effects are mostly in the range -0.02 to -0.045. The latter value corresponds to approximately 31 percent of the outcome's standard deviation and suggests that the financial circumstances before payday led to 0.045 fewer correct answers per second in the Stroop task. For the older individuals with a lower current income, the effect estimates are between -0.01 and -0.02. Similar to Panel A, Panel B shows that the causal forest using the number of correct answers as the outcome also estimates particularly detrimental effects for individuals with a current income of at most around \$750, and who are either younger or older. For the older individuals, the especially detrimental effect estimates are again concentrated in the approximate age range 70 to 80 years. However, they now actually also exceed the \$750 threshold. The most detrimental effect estimate in the Panel B heatmap equals -1.47, which corresponds to approximately 14 percent of the standard deviation of the outcome. Panel C displays the estimated effects for the outcome total response time. Similar to the other two panels, the heatmap also yields detrimental effects for individuals whose current income is below \$750, and among the lower current income individuals, the causal forest again estimates particularly detrimental effects for younger individuals (up to around 27 years) and older individuals (approximately above age 67). The estimated effects in the most detrimental category are located at the ages 78 to 82 years for current income levels of up to \$425, and then at the ages between around 70 and 83 years for current income between approximately \$425 and \$750. In this category, the causal forest gives effect estimates on the total response time of up to 4.18 seconds, or 18 percent of the outcome's standard deviation. Thus, the

[15]See Appendix D for further details.

heatmaps in Figure 2 suggest that the poorer financial circumstances before payday impede, in particular, the cognition of younger and older individuals with a lower current income. The negative effect on the number of correct answers per second appears to result not only from fewer correct answers given but also a slower total response time.

To gain a deeper understanding of the detrimental effects of the financial circumstances before payday, we next zoom in on two regions in which the heatmaps indicate particularly harmful effects. Specifically, we estimate effects for a typical younger individual aged 20 and a typical older individual aged 75, who both have a current income of $450. We refer to these individuals as typical because we set all other 35 covariates for estimating the effects according to the characteristics in a neighborhood of a given age-current income combination: that is, we construct a five-year age and $250 current–income window centered at the respective age–current income combination and determine the covariate values within this window using the same procedure as for creating the heatmaps above.[16] The first row in the panels of Figures 3 and 4 gives the estimates for the two typical individuals and all three outcomes. We call these estimates the typical individual baseline estimates. To study how changing the other 35 covariates affects the effect estimates, the panels then show, in the rows below the first row, estimates for which we change one characteristic of a given typical individual at a time, leaving all other variables constant. The empty rows indicate how the covariates are set for a typical individual. For example, for the younger individual in Figure 3, the row labeled 'Male = 0' is empty. This indicates that the younger typical individual is female. The row labeled 'Male = 1' then shows the effect estimate when we change the typical individual's gender from female to male. Similarly, the row labeled 'unemployed = 1' gives the effect estimate when we change the individual's employment status from working to unemployed (every time leaving all other covariates unchanged). In both figures, the horizontal bars indicate 90 percent confidence intervals.

The first row of Panel A in Figures 3 and 4 shows that the causal forest estimates a negative effect of the financial circumstances before payday on the number of correct answers per second for the younger and older typical individuals. For the younger individual, the estimated effect is -0.0477, and significant at the 1 percent level. The effect size corresponds to approximately one third of the outcome's standard deviation. For the older individual, the effect estimate is -0.0370, or approximately 26 percent of the standard deviation of the

[16] See Appendix D for further details.

outcome. The estimate is significant at the 5 percent level.[17]

In line with the findings from the heatmaps, row one in Panels B and C in Figures 3 and 4 suggests that the detrimental effect on the main outcome results from the financial circumstances before payday having a detrimental effect on both its numerator and denominator. The estimate for the effect on the number of correct answers is negative, and the estimate for the effect on the total response time is positive. However, only the response-time effect estimate for the older individual is significant at a conventional level.

The rows below the first row in Figures 3 and 4 show that changing a single characteristic of the two typical individuals does not yield estimates that differ much compared with the baseline estimates. The sign of the effect estimates never changes, and the magnitude of the point estimates remains similar.[18] This behavior is in line with the conclusion from the variable importance plots that age and current income are by far the most important variables for effect heterogeneity.

4.3 Subgroup Analysis

Overall, our heterogeneity analysis suggests that the poorer financial circumstances before payday are especially detrimental for individuals who have a current income below approximately \$750 and whose age is either roughly below 30 or above 70 years. Based on this insight, we next estimate average treatment effects for this subgroup of interest in our sample. Subsequently, to verify the findings in our main analysis sample, we estimate average treatment effects for the subgroup of interest in an independent experiment that Carvalho et al. (2016) conducted in their second online panel. Based on this additional experiment, we only perform a traditional OLS subgroup analysis. We do not use the additional experiment in our heterogeneity analysis using the causal forest method.

To estimate average effects in our main analysis sample, we use the augmented inverse propensity weighted estimator (Robins and Rotnitzky 1995) implemented in the grf R package. The estimator uses the causal forest estimates for all individuals in the subgroup of interest to form the average effect estimates. Table 4 presents the estimation results for the subgroup analysis. Column (1) shows the estimate for the effect of the financial circumstances before payday on the number of correct answers per second. The estimation yields

[17] Appendix Tables 4 and 5 display the estimates discussed in the text.

[18] Appendix E shows that the conclusions based on other typical individuals in the vicinity of the two typical individuals discussed in the text are the same.

an effect estimate of -0.098, which corresponds to approximately 69 percent of the outcome's standard deviation. The estimate is significant at the 1 percent level. Columns (2)–(3) display the results from the estimations that use the numerator and denominator of our main outcome as dependent variables. Both regressions also give harmful effect estimates, which are significant at least at the 5 percent level. Thus, in line with the findings from our heterogeneity analysis, the estimations yield detrimental effects on cognition of the poorer financial circumstances before payday for our subgroup of interest.

To verify the validity of this finding, we next estimate treatment effects for our subgroup of interest in the Flanker task experiment of Carvalho et al. (2016). The authors conducted this experiment in the second online panel that they used, the RAND American Life Panel. Just as the Stroop task studied in our main analysis, the Flanker task measures cognitive control, and its goal is also to give correct answers to a repeated stimulus as quickly as possible. Carvalho et al. (2016) ran the experiment with 20 trials per participant.

Panel A in Table 5 replicates Carvalho et al.'s (2016) OLS estimates for the Flanker task. The regressions do not suggest that the poorer financial circumstances before payday have an effect on cognition in the full sample. The estimated effect on the probability of giving a correct answer in a trial, in Column (1), and on the (log) time that an individual took to respond to a trial, in Column (2), is close to zero and insignificant at the 10 percent level. Panel B displays the analogous estimates for our subgroup of interest. While the estimate in Column (1) does not suggest there to be an effect on the probability of giving a correct answer, the estimate in Column (2) does indeed suggest a detrimental effect on the log response time per trial. The latter estimate is 0.274 and significant at the 1 percent level. This suggests that the individuals responded on average approximately 27% more slowly to the trials of the Flanker task due to the poorer financial circumstances before payday. Thus, in line with the results of our main analysis based on the KnowledgePanel, the analysis based on the American Life Panel also yields detrimental cognitive effects of the poorer financial circumstances before payday for younger and older individuals who have a lower income around the time of the experiment.

5 Conclusion

In this paper, we examine heterogeneity in the effect of financial circumstances on cognition. Our analysis is based on data from an experiment by Carvalho et al. (2016), which randomly assigned low-income individuals in the US to perform a cognitive test before or after payday. To explore heterogeneity in the effect of the poorer financial circumstances before payday, we use the causal forest method by Athey et al. (2019), which is designed for studying heterogeneous treatment effects in experiments.

The results of our analysis suggest that financial circumstances have heterogeneous effects on cognition. While in our full sample the estimations do not suggest that the poorer financial circumstances before payday affect cognition, we do find detrimental effects for younger and older individuals who received a very low income around the time of the experiment. Specifically, our findings suggest that cognitive test performance was worse among those who received an income of less than $750 at the time of the experiment and whose age was below 30 or above 70 years. We also find detrimental cognitive effects for this group of individuals in an additional, independent, experiment conducted by Carvalho et al. (2016), which we do not use in our heterogeneity analysis. Among the 37 covariates included in our analysis, age and current income appear to be by far the most important for effect heterogeneity. All of the other covariates, such as marital status, household size and education, do not appear to play an important role.

We derive a number of policy recommendations from our findings: First, to address the potential negative cognitive effects of poor financial circumstances, it may be especially beneficial when designing poverty reduction measures to target these at individuals with very few current financial resources and who are either relatively young or old. Second, for this group of individuals, it may prove helpful for public policy to take into account a possible variation in cognitive capacity over payment cycles. For example, to prevent potentially poor decision making due to limited cognition, public administration could try to avoid scheduling appointments with the affected individuals at the end of their payment cycles. Because the payment cycles of welfare programs, such as the food stamp program, are generally regular and set far in advance, this appears to be a feasible option, especially in cases where individuals receive welfare payments.

A fruitful avenue for further research might be to explore why the financial circumstances before payday had detrimental effects for some, but not all, individuals in the experiment.

A low current income, for example, may capture particularly poor financial circumstances before payday, and younger and older individuals may be especially worried about these. To gain a deeper understanding of the mechanisms at play, it would be helpful to obtain a larger experimental data set, which focuses on our identified subgroup of affected individuals and would allow for a more detailed analysis.

References

Athey, S., and Imbens, G. (2016). Recursive Partitioning for Heterogeneous Causal Effects. *Proceedings of the National Academy of Sciences*, *113*(27), 7353–7360.

Athey, S., Tibshirani, J., and Wager, S. (2019). Generalized Random Forests. *Annals of Statistics*, *47*(2), 1148–1178.

Bertrand, M., Mullainathan, S., and Shafir, E. (2004). A Behavioral-Economics View of Poverty. *American Economic Review*, *94*(2), 419–423.

Bertrand, M., Mullainathan, S., and Shafir, E. (2006). Behavioral Economics and Marketing in Aid of Decision Making among the Poor. *Journal of Public Policy and Marketing*, *25*(1), 8–23.

Bourke, N., Horowitz, A., and Roche, T. (2012). *Payday Lending in America: Who Borrows, Where They Borrow, and Why*, Pew Charitable Trusts. Payday Lending in America Series.

Breiman, L. (2001). Random Forests. *Machine Learning*, *45*(1).

Carvalho, L. S., Meier, S., and Wang, S. W. (2016). Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday. *American Economic Review*, *106*(2), 260–284.

Clotfelter, C. T., Cook, P. J., Edell, J. A., and Moore, M. (1999). *State Lotteries at the Turn of the Century: Report to the National Gambling Impact Study Commission*. Retrieved August 30, 2018, from http://govinfo.library.unt.edu/ngisc/reports/lotfinal.pdf

Dube, S. R., Asman, K., Malarcher, A., and Carabollo, R. (2009). *Cigarette Smoking among Adults and Trends in Smoking Cessation – United States, 2008*, Centers for Disease Control and Prevention. Morbidity and Mortality Weekly Report Series, 58(44), 1227–1232.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning*. New York: Springer.

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

Mastrobuoni, G., and Weinberg, M. (2009). Heterogeneity in Intra-monthly Consumption Patterns, Self-Control, and Savings at Retirment. *American Economic Journal: Economic Policy*, *1*(2), 163–189.

Robins, J. M., and Rotnitzky, A. (1995). Semiparametric Efficiency in Multivariate Regression Models with Missing Data. *Journal of the American Statistical Association*, *90*(429), 122–129.

Ross, J. S., Bernheim, S. M., Bradley, E. H., Teng, H.-M., and Gallo, W. T. (2007). Use of Preventive Care by the Working Poor in the United States. *Preventive Medicine*, *44*(3), 254–259.

Shapiro, J. M. (2005). Is there A Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle. *Journal of Public Economics*, *89*(2-3), 303–325.

Stephens Jr., M. (2003). ”3rd of tha Month”: Do Social Security Recipients Smooth Consumption between Checks? *American Economic Review*, *93*(1), 406–422.

Stephens Jr., M. (2006). Paycheque Receipt and the Timing of Consumption. *Economic Journal*, *116*(513), 680–701.

Tibshirani, J., Athey, S., Wager, S., Friedberg, R., Miner, L., and Wright, M. (2018). *grf: Generalized Random Forests (Beta). R package version 0.9.6.*

Wager, S., and Athey, S. (2018). Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests. *Journal of the American Statistical Association*, *113*(523), 1228–1242.

Tables

Table 1. Definitions and Descriptive Statistics for the Outcomes and Regressor

	Definition	Mean	Standard Deviation
Outcomes			
Correct answers per second	Number of correct answers that individual i gave across all 48 Stroop task trials divided by the total time in seconds that it took i to complete the entire Stroop task.	0.446	0.143
Number of correct answers	Number of correct answers that individual i gave across all 48 trials of the Stroop task.	42.899	10.565
Total response time in seconds	Total time in seconds that it took individual i to complete the entire Stroop task.	100.476	22.816
Regressor of interest			
Before payday	=1 if individual i was assigned to be surveyed before payday.	0.509	0.500

Notes: $N = 2,480$. The data are from the KnowledgePanel experiment by Carvalho et al. (2016).

Table 2. Descriptive Statistics for the Covariates

	Mean	Standard Deviation
Age	55.947	17.423
Male	0.334	0.472
Household size	1.944	1.192
Household head	0.846	0.361
Children in household	0.167	0.373
Metropolitan area	0.804	0.397
Current income	1737.987	1321.136
Share of payday pay amount relative to current income	0.762	0.278
Financial strain		
Live from paycheck to paycheck	0.489	0.500
Caloric crunch	0.470	0.499
Liquidity constrained	0.503	0.500
Financial hardship	0.414	0.493
Marital status		
Married	0.335	0.472
Divorced	0.276	0.447
Widowed	0.139	0.346
Never married	0.250	0.433
Race		
White	0.761	0.426
Black	0.100	0.300
Hispanic	0.082	0.274
Other race	0.057	0.232
Employment status		
Working	0.287	0.452
Unemployed	0.063	0.244
Disabled	0.199	0.399
Retired	0.388	0.487
Other employment status	0.062	0.242
Education		
Less than high school	0.063	0.244
High school	0.254	0.435
Some college	0.417	0.493
College	0.266	0.442
Annual household income		
Less than $5,000	0.048	0.215
Between $5,000 and $10,000	0.100	0.300
Between $10,000 and $15,000	0.143	0.350
Between $15,000 and $20,000	0.120	0.325
Between $20,000 and $25,000	0.149	0.356
Between $25,000 and $30,000	0.143	0.350
Between $30,000 and $35,000	0.140	0.347
Between $35,000 and $40,000	0.156	0.363

Notes: $N = 2,480$. The data are from the KnowledgePanel experiment by Carvalho et al. (2016). The dummy category other race also includes individuals of mixed ethnicity; unemployed also includes temporarily laid off individuals, and working also includes self employed individuals. In the order of the four financial strain variables listed, each respective dummy equals one if an individual i) agrees or strongly agrees with the statement 'I live from paycheck to paycheck', ii) had to reduce consumption at the end of a pay cycle, iii) could not, or would have to do something drastic to, raise $2,000 in one week for an emergency, iv) experienced at least one out of ten hardships related to not having enough money in the past 12 months. For the ten hardships, see Table C4 in the online appendix of Carvalho et al. (2016).

Table 3. OLS Average Effect Estimates

Outcome	Correct answers per second	Number of correct answers	Total response time (in seconds)
	(1)	(2)	(3)
Before payday	0.007	0.183	-1.062
	(0.006)	(0.425)	(0.916)
Constant	0.443^{***}	42.805^{***}	101.017^{***}
	(0.004)	(0.305)	(0.643)

Notes: $N = 2,480$. The data are from the KnowledgePanel experiment by Carvalho et al. (2016). Heteroskedasticity-robust standard errors are in parentheses.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Table 4. Subgroup Average Effect Estimates

Outcome	Correct answers per second	Number of correct answers	Total response time (in seconds)
	(1)	(2)	(3)
Before payday	-0.098^{***}	-3.660^{**}	11.823^{***}
	(0.023)	(1.539)	(2.890)

Notes: N=117. The data are from the KnowledgePanel experiment by Carvalho et al. (2016). Standard errors are in parentheses. The sample includes all individuals who have a current income below $750 and whose age is either below 30 or above 70 years. The estimates are obtained via an augmented inverse propensity weighted estimator, which is based on the causal forest estimates for the individuals in the sample. For more information, see Section 4.3.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Table 5. Average Effect Estimates in an Independent Experiment

Outcome	Correct answer	Log response time per trial
	(1)	(2)
Panel A. Full sample		
Before payday	0.007	0.016
	(0.010)	(0.028)
Constant	0.863***	8.060***
	(0.012)	(0.030)
N	20,557	20,557
Individuals	1,076	1,076
Panel B. Subgroup: Current income below \$750 and age below 30 or above 70 years		
Before payday	0.045	0.274***
	(0.041)	(0.099)
Constant	0.845***	7.908***
	(0.047)	(0.107)
N	1,590	1,590
Individuals	85	85

Notes: The data are from the Flanker task experiment in the RAND American Life Panel by Carvalho et al. (2016). The table reports OLS estimates. Standard errors clustered at the individual level are in parentheses. The regressions include trial-specific dummies. The outcome correct answer is a dummy which equals one if individual i answered a trial correctly. The outcome log response time per trial measures the log time in milliseconds that individual i took to respond to a trial. Panel A replicates the results in Carvalho et al.'s (2016) Table 6.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Figures

Figure 1. Variable Importance Plots for the Causal Forests

Panel A. Correct Answers per Second

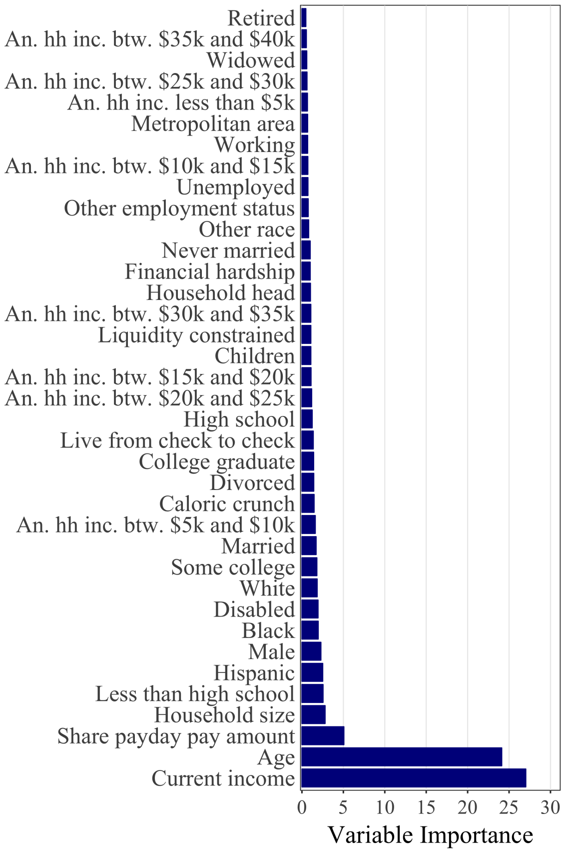

Panel B. Number of Correct Answers

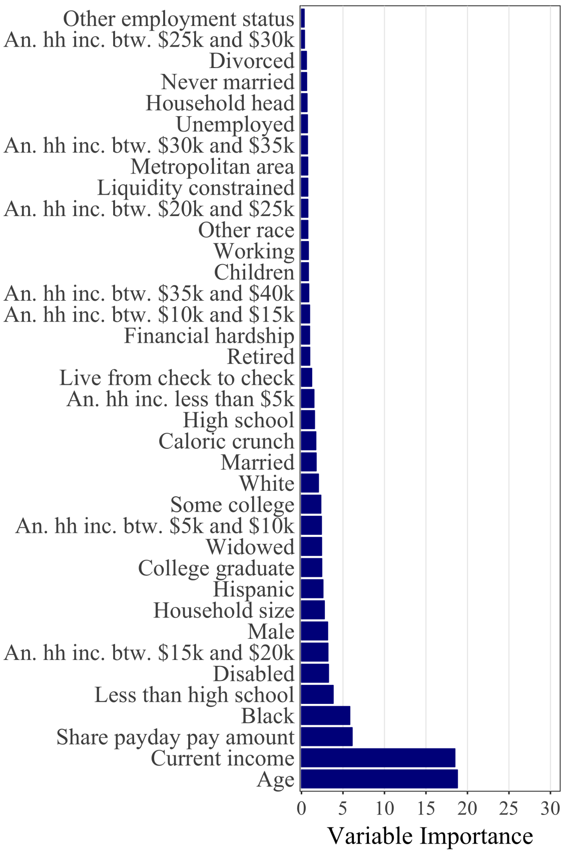

Panel C. Total Response Time

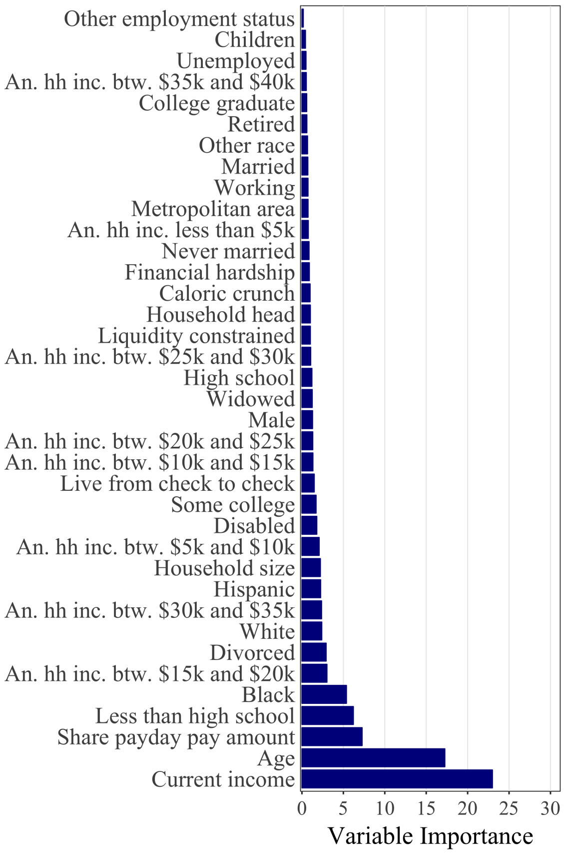

N=2,480. The data are from the KnowledgePanel experiment by Carvalho et al. (2016). The variable importance measures are calculated based on the causal forest estimations. For more information, see Section 4.2.

Figure 2. Causal Forest Effect Estimates over an Age-Current Income Grid

Panel A. Correct Answers per Second

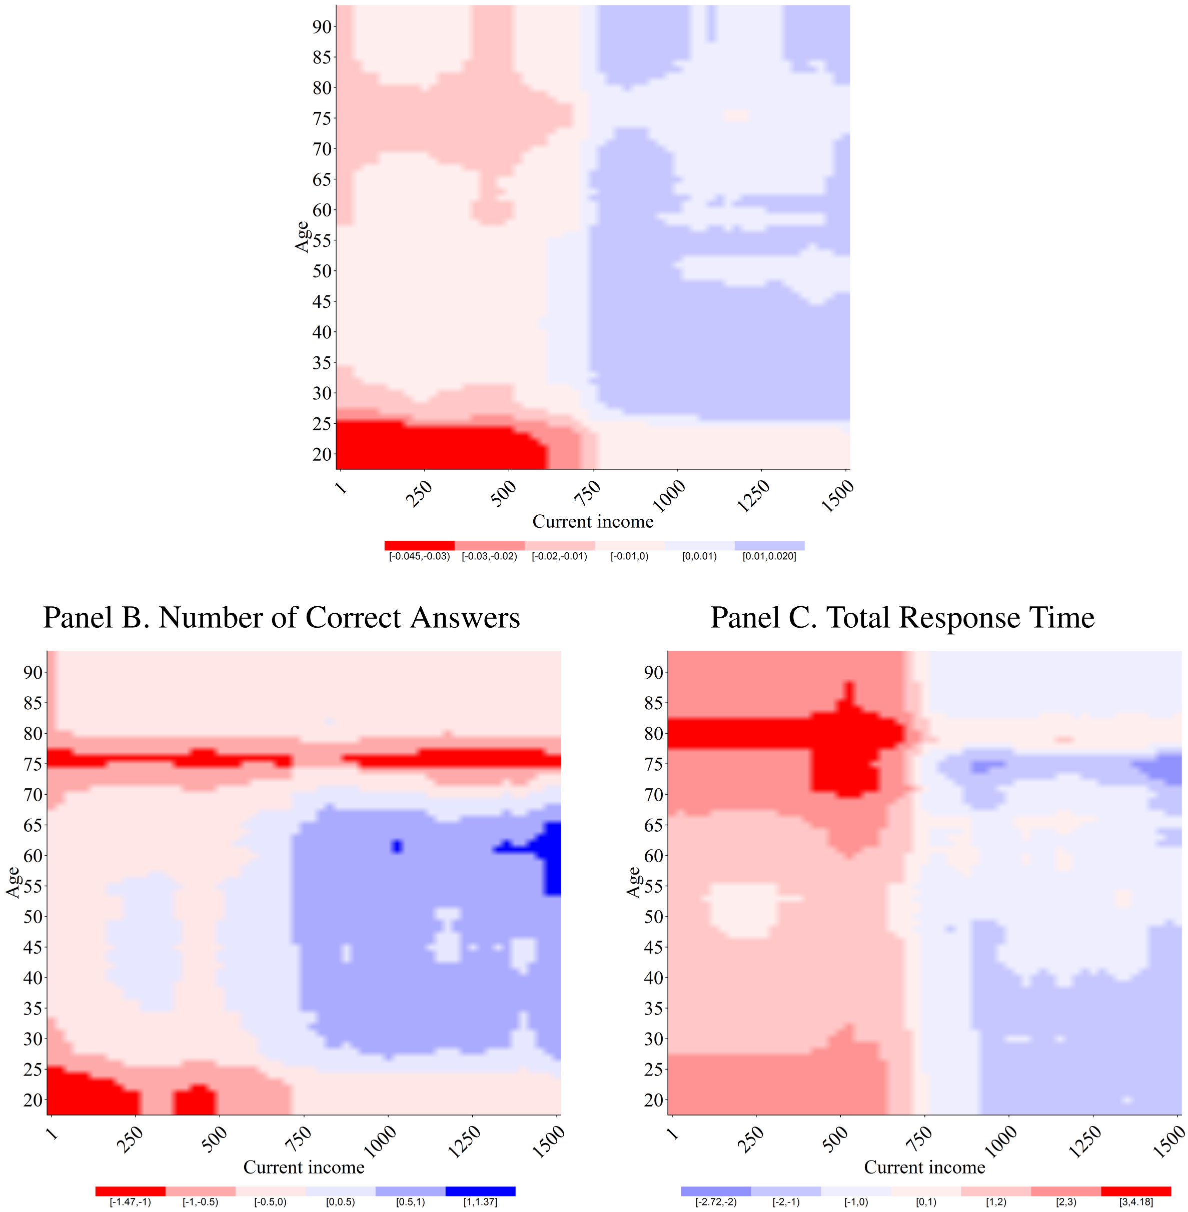

N=2,480. The data are from the KnowledgePanel experiment by Carvalho et al. (2016). The heatmaps show causal forest estimates for the effect of the poorer financial circumstances before payday on our three cognition outcomes. For more information, see Section 4.2.

Figure 3. Causal Forest Effect Estimates for the Typical Younger Individual
Age = 20, current income = $450

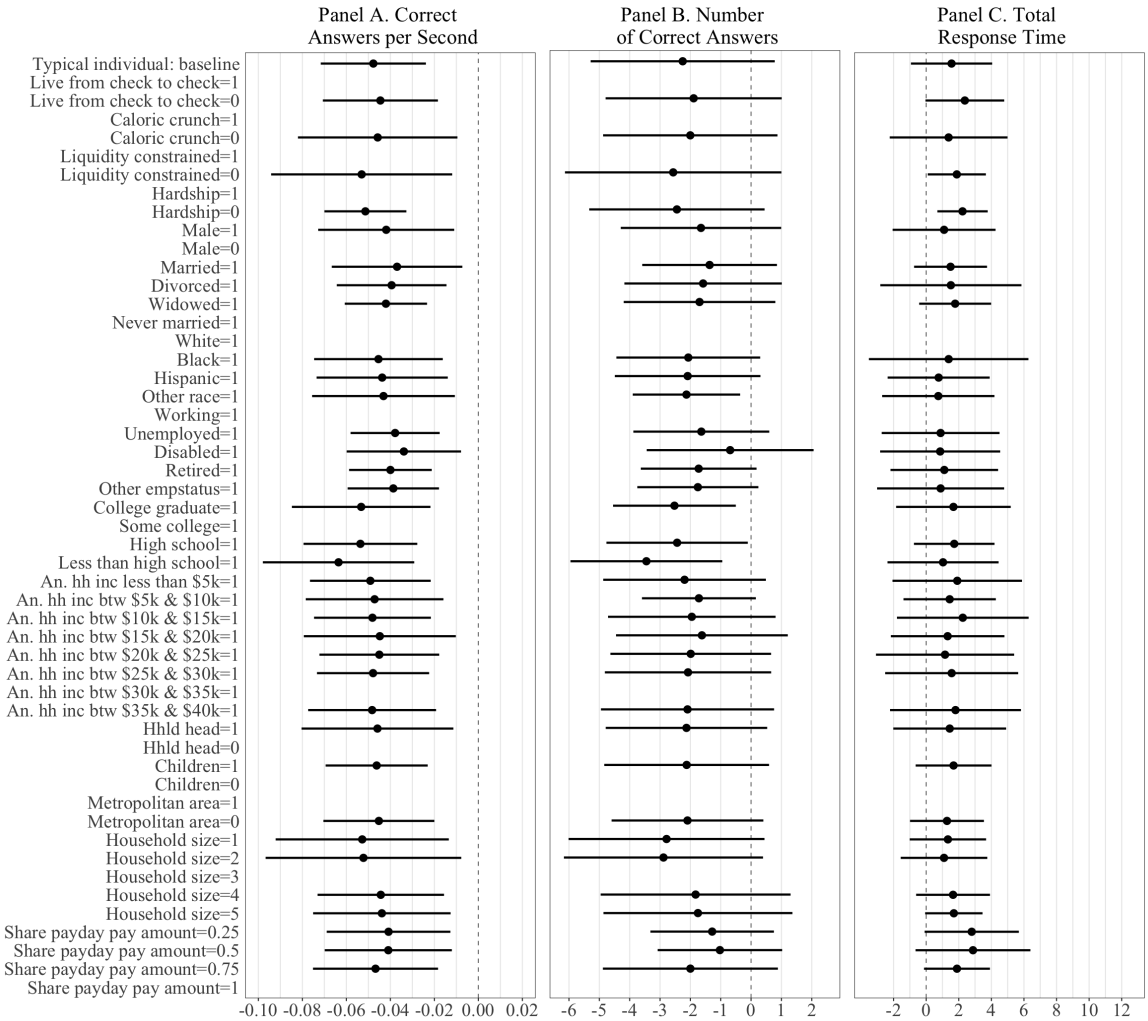

N=2,480. The data are from the KnowledgePanel experiment by Carvalho et al. (2016). The plots show causal forest estimates for the effect of the poorer financial circumstances before payday on our three cognition outcomes. The horizontal bars indicate 90 percent confidence intervals. For the covariates household size and share payday pay amount, the plots give effect estimates at selected points. For more information, see Section 4.2.

Figure 4. Causal Forest Effect Estimates for the Typical Older Individual

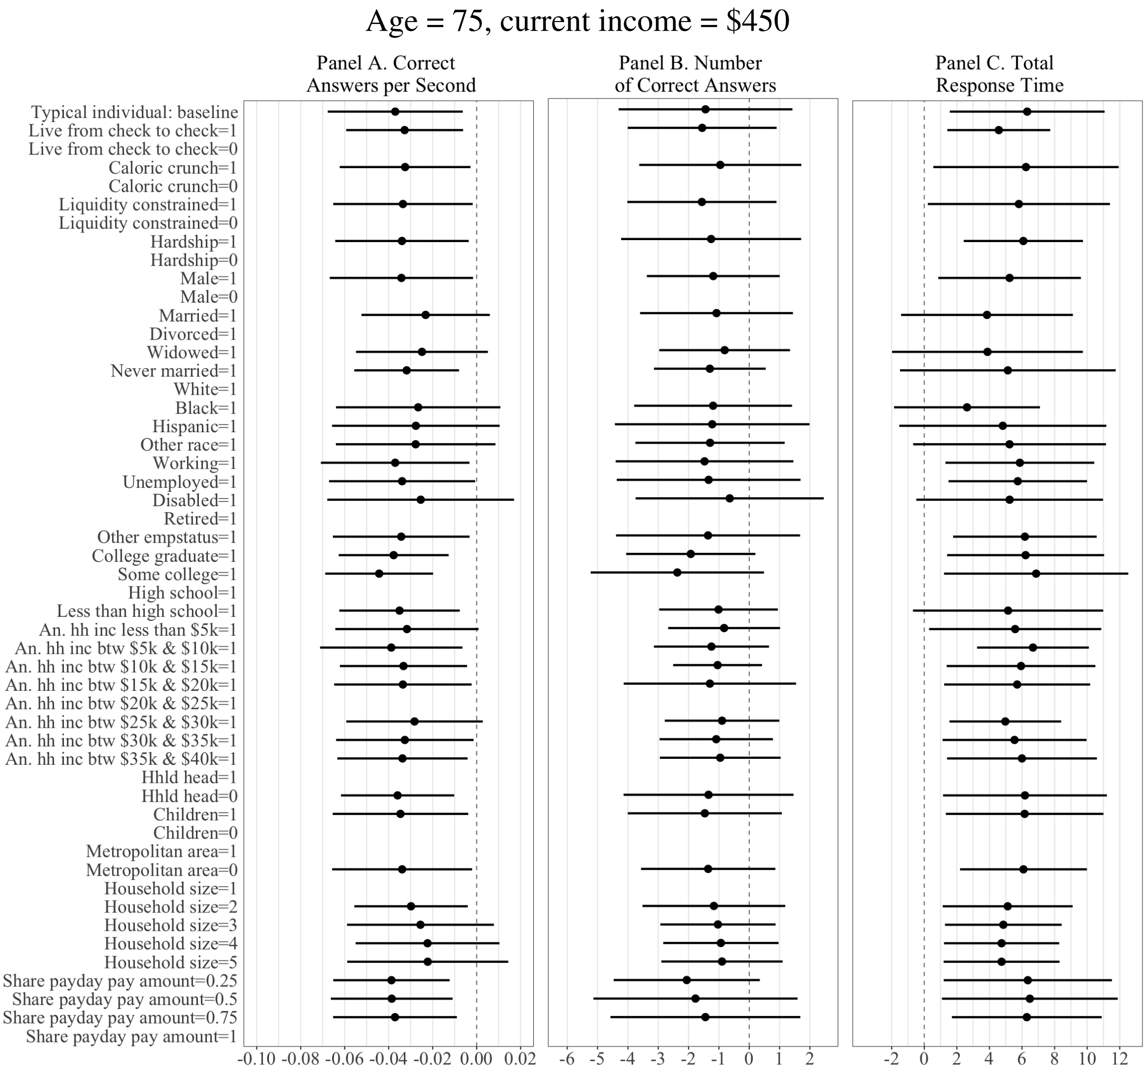

N=2,480. The data are from the KnowledgePanel experiment by Carvalho et al. (2016). The plots show causal forest estimates for the effect of the poorer financial circumstances before payday on our three cognition outcomes. The horizontal bars indicate 90 percent confidence intervals. For the covariates household size and share payday pay amount, the plots give effect estimates at selected points. For more information, see Section 4.2.

Appendix

A Derivation of the Causal Forest Estimator

The causal forest estimator $\hat{\tau}(x)$ for $\tau(x)$ in the random effects model posited in Section 3 is based on the two local moment equations

$$E\left(Y_i - \tau(x)D_i - c(x)|X_i = x\right) = 0 \tag{4}$$

$$E\left(\left(Y_i - \tau(x)D_i - c(x)\right)D_i|X_i = x\right) = 0, \tag{5}$$

where $c(x) = E(\epsilon_i|X_i = x)$ is an intercept term. All other quantities are defined as in the main text. The estimator $\hat{\tau}(x)$ is now obtained by minimizing an empirical version of the two local moment equations:

$$(\hat{\tau}(x), \hat{c}(x)) = argmin_{\tau(x),c(x)} \left\| \sum_{i=1}^{n} \alpha_i(x) \begin{pmatrix} Y_i - \tau(x)D_i - c(x) \\ (Y_i - \tau(x)D_i - c(x))D_i) \end{pmatrix} \right\|_2 . \tag{6}$$

The resulting causal forest estimator can be written as

$$\hat{\tau}(x) = \frac{\sum_{i=1}^{n} \alpha_i(x)(Y_i - \bar{Y}_\alpha)(D_i - \bar{D}_\alpha)}{\sum_{i=1}^{n} \alpha_i(x)(D_i - \bar{D}_\alpha)^2}, \tag{7}$$

where $\bar{Y}_\alpha = \sum_{i=1}^{n} \alpha_i(x)Y_i$, $\bar{D}_\alpha = \sum_{i=1}^{n} \alpha_i(x)D_i$, and $\alpha_i(x)$ are the similarity weights. It holds that $\sum_{i=1}^{n} \alpha_i(x) = 1$.

Equation (7) is the expression for the causal forest estimator in Section 6 of Athey et al. (2019). To obtain the formulation of the estimator in Equation (2) of our main text, rewrite Equation (7) as follows. For the numerator, we have

$$\begin{aligned}
&\sum_{i=1}^{n} \alpha_i(x)(Y_i - \bar{Y}_\alpha)(D_i - \bar{D}_\alpha) \\
&= \sum_{i=1}^{n} \alpha_i(x)Y_iD_i - \left(\sum_{i=1}^{n} \alpha_i(x)D_i\right)\left(\sum_{i=1}^{n} \alpha_i(x)Y_i\right)
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{\{i:D_i=1\}} \alpha_i(x)Y_i - \left(\sum_{\{i:D_i=1\}} \alpha_i(x)\right)\left(\sum_{\{i:D_i=1\}} \alpha_i(x)Y_i + \sum_{\{i:D_i=0\}} \alpha_i(x)Y_i\right) \\
&= \left(1 - \sum_{\{i:D_i=1\}} \alpha_i(x)\right)\sum_{\{i:D_i=1\}} \alpha_i(x)Y_i - \left(\sum_{\{i:D_i=1\}} \alpha_i(x)\right)\sum_{\{i:D_i=0\}} \alpha_i(x)Y_i \\
&= \left(\sum_{\{i:D_i=0\}} \alpha_i(x)\right)\sum_{\{i:D_i=1\}} \alpha_i(x)Y_i - \left(\sum_{\{i:D_i=1\}} \alpha_i(x)\right)\sum_{\{i:D_i=0\}} \alpha_i(x)Y_i \qquad (8)
\end{aligned}
$$

For the denominator, we have

$$
\begin{aligned}
\sum_{i=1}^{n} \alpha_i(x)(D_i - \bar{D}_\alpha)^2 &= \sum_{i=1}^{n} \alpha_i(x)D_i - \left(\sum_{i=1}^{n} \alpha_i(x)D_i\right)^2 \\
&= \left(\sum_{\{i:D_i=1\}} \alpha_i(x)\right)\left(1 - \sum_{\{i:D_i=1\}} \alpha_i(x)\right) \\
&= \left(\sum_{\{i:D_i=1\}} \alpha_i(x)\right)\left(\sum_{\{i:D_i=0\}} \alpha_i(x)\right) \qquad (9)
\end{aligned}
$$

The derivations for the numerator and denominator exploit $\sum_{i=1}^{n} \alpha_i(x) = 1$ and $D_i^2 = D_i$. Plugging expression (8) for the numerator and expression (9) for the denominator into $\hat{\tau}(x)$ from (7) yields

$$
\begin{aligned}
\hat{\tau}(x) &= \frac{\left(\sum_{\{i:D_i=0\}} \alpha_i(x)\right)\sum_{\{i:D_i=1\}} \alpha_i(x)Y_i - \left(\sum_{\{i:D_i=1\}} \alpha_i(x)\right)\sum_{\{i:D_i=0\}} \alpha_i(x)Y_i}{\left(\sum_{\{i:D_i=1\}} \alpha_i(x)\right)\left(\sum_{\{i:D_i=0\}} \alpha_i(x)\right)} \\
&= \sum_{\{i:D_i=1\}} \frac{\alpha_i(x)}{\sum_{\{i:D_i=1\}} \alpha_i(x)} Y_i - \sum_{\{i:D_i=0\}} \frac{\alpha_i(x)}{\sum_{\{i:D_i=0\}} \alpha_i(x)} Y_i, \qquad (10)
\end{aligned}
$$

which is the expression for the causal forest estimator in the main text.

B Tree Example

The figure below this paragraph shows an example of a single small (regression) tree. The tree is built on a sample of size $n = 50$. The data used to construct the tree includes the continuous covariates $X1$, $X2$ and the continuous outcome Y. In the first step, starting from the top of the figure, the tree splits the full sample into two partitions based on the variable $X1$. All observations with an $X1 < 8.6$ are put into the 'left' partition and all observations with an $X1 \geq 8.6$ are put into the 'right' partition. Analogously, the tree subsequently splits the resulting 'left' partition on variable $X2$ and the resulting 'right' partition on variable $X1$ again. The splitting procedure yields four leaves, which are shown at the bottom of the figure. For each leaf, the tree calculates the average outcome Y by averaging over all Y values of all observations that fall into the respective leaf. The averages are then used for predicting Y. For example, for an observation with an $X1 < 8.6$ and $X2 \geq 3.5$, the tree predicts an outcome value of 10.

Regression Tree Example

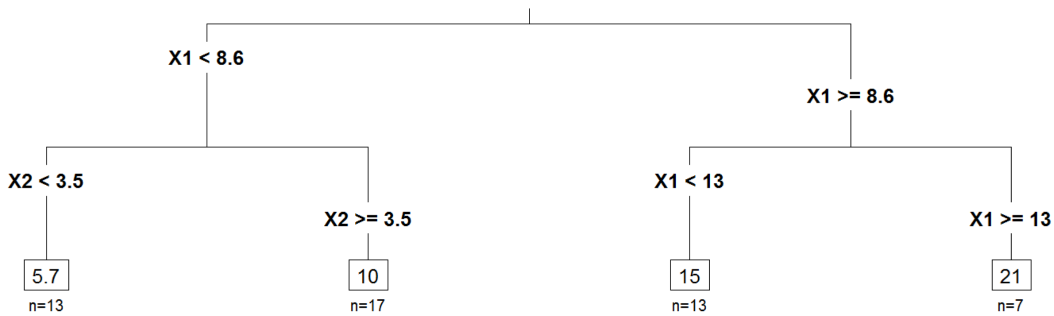

Notes: The values in the boxes correspond to the average outcome Y over all observations that fall into a respective leaf. The number of observations within each leaf is denoted by n.

C Calculation of the Variable Importance Measure

The variable importance measure that we use in our causal forest analysis is implemented in the function variable_importance in the R package grf. We multiply the measure by 100 for readability. The function requires to set the maximum tree depth up to which the measure considers splits, and a decay exponent that controls how the weight that the splits receive in the overall measure changes as the tree depth increases.[19] We use the default values of the

[19]For a given tree, the split at depth one corresponds to the first split that a tree places, starting from the entire subsample, and splitting it into two partitions. The splits at depth two then correspond to the splits that

variable_importance function for the two parameters: we set the maximum tree depth to four and the decay exponent to two. For variable X_k, the measure is calculated as follows:

$$vi(X_k) = \left(\sum_{j=1}^{4} w_j \frac{n_{jk}}{n_j} \right) \times 100, \tag{11}$$

where n_{jk} is the number of times that all of the trees of the causal forest together split on variable X_k at tree depth $j, j = 1, \ldots, 4$. n_j is the number of times that the trees split at depth j, and $w_j = \frac{j^{-2}}{\sum_{l=1}^{4} l^{-2}}$ is a tree depth-specific weight that determines the importance of splits at a given depth.

In short, the variable importance measure $vi(X_k)$ is a weighted sum of the relative splitting frequencies for X_k over the depths $j = 1, \ldots, 4$, where the weight of the relative splitting frequencies decreases as the tree depth increases.

D Procedure to Set the Covariates

For creating the heatmaps in Figure 2, we set the covariates household size and share of payday pay amount relative to current income to their full sample median values. All other covariates, which are dummies, we set according to the most frequently occurring characteristics in the full sample. To give two more examples in addition to the example in Section 4.2, Table 2 shows that the most frequent marital status category, with 33.5 percent, is married. Thus, we set the dummy married equal to one and all other marital status dummies to zero. Furthermore, Table 2 shows that 80.4 percent of individuals live in a metropolitan area. Accordingly, we set the dummy metropolitan area equal to one.

To obtain the estimates for the typical individuals in Appendix Tables 4 and 5 and the first row of Figures 3 and 4, we proceed analogously to the covariate setting procedure for the heatmaps. However, rather than setting the variables according to the full sample characteristics, we determine the covariate values according to the characteristics in a five–year age and \$250 current–income window, which is centered at the age–current income combination for which we want to estimate an effect. For example, for the typical older individual in Figure 4, the relevant window for setting the covariates ranges from 73 to 77 years of age and \$325.5 to \$574.5 of current income. If there are tie categories in categorical variables

the tree performs starting from the two partitions created at depth 1. The next depths follow analogously.

or dummies that refer to an ordinal characteristic, such as income or education, we select the lowest tie category.[20] For example, if there are equally many individuals in a respective age–current income window with a high school degree and some college, we set the dummy high school to one and all other education dummies to zero. If there are tie categories in non-ordinal characteristics, such as martial status or being liquidity constrained, we set the respective covariates by extending the age–current income window by one year and \$100 in each direction, i.e., we use a seven–year age and \$450 current–income window.[21]

E Estimates in the Vicinity of the Two Typical Individuals

In our main analysis, we estimate effects for two typical individuals who have a current income of \$450 and whose age is 20 and 75 years, respectively. To assert that the insights based on the two typical individuals are not sensitive to the specific choice of the age–current income combination, we additionally estimate effects for other typical individuals that are in the vicinity of our two typical individuals from the main analysis, where the heatmaps also indicate pronounced detrimental effects. Specifically, we increase and decrease, respectively, age by one and two years and current income by \$25 and \$50 relative to the typical individuals from the heterogeneity analysis. We estimate the effects analogously to the typical individual baseline estimates in Figures 3 and 4. Appendix Tables 4 and 5 present the estimates for the other typical younger and older individuals. In the interest of space, we do not display the effect estimates when varying the other 35 covariates. However, very similar to the findings in our main analysis, varying the other covariates one by one does also not change the estimates much relative to the baseline estimates.

The comparison between the estimates for the two typical individuals from our main analysis, shown in the gray shaded areas of Appendix Tables 4 and 5, and the other typical individuals shows that overall the vicinity estimates are quite similar to the main analysis estimates. For our main outcome, correct answers per second, Panel A in both tables shows that the estimates for the other typical individuals are also always negative and of a similar

[20]Similarly, if the median household size, as determined by R, is a non-integer value, we set the household size to the largest integer below the respective median household size. For example, a median household size of 3.5, we set to 3.

[21]For the typical individual with age 20 and current income equal to \$400 in Appendix Table 3, extending the age–current income window does not break the tie in the variable household head. In this case, we set household head equal to zero. Setting household head equal to one instead does not change the conclusions for the corresponding estimate.

magnitude as for the respective younger or older typical individual from the main analysis. For the typical younger individuals, all estimates, except for one, are significant at conventional levels. For the typical older individuals, the estimates sometimes lose significance at the 10 percent level.

Panel B in Appendix Tables 4 and 5 shows that the estimates for the outcome number of correct answers are also always negative and the point estimates appear quite similar to the respective estimate for the main analysis typical individual, considering the magnitude of the standard errors. As in the main analysis, the estimates are insignificant at the 10 percent level in most regressions. Similar to the findings in Panel B, the estimates for the other typical individuals using the outcome total response time in Panel C are also not substantially different from the respective estimate for the typical individual in the main analysis. In all regressions, the estimations yield positive effect estimates that are insignificant at the 10 percent level for the younger individuals, and mostly significant at conventional levels for the older individuals.

F Appendix Tables and Figures

Appendix Table 1. Variation in Financial Resources at Payday

Outcome	Cash	Checking and savings	Total expenditures
	(1)	(2)	(3)
Panel A. OLS regressions			
Before payday	-33.39	-6032.75	-542.88
	(73.72)	(5083.40)	(378.69)
Constant	273.18^{***}	15520.66^{***}	1279.50^{***}
	(55.52)	(5000.96)	(371.45)
Panel B. Median regressions			
Before payday	-5.00	-500.00^{***}	-200.00^{***}
	(4.26)	(122.05)	(33.03)
Constant	50.00^{***}	1500.00^{***}	600.00^{***}
	(2.19)	(109.85)	(26.08)
Panel C. p-values for Wilcoxon tests of equality of distributions			
	0.01	0.00	0.00
N	2,295	2,127	2,296

Notes: The data are from the KnowledgePanel experiment by Carvalho et al. (2016). For the OLS regressions, heteroskedasticity-robust standard errors are in parentheses. For the median regressions, bootstrap standard errors based on 1,000 replications are in parentheses. Compared with the analogous results in Carvalho et al.'s (2016) Table 1, only the before payday estimate in the OLS regression using the outcome total expenditures and the before payday estimate in the median regression using the outcome cash loses significance in our sample, which is smaller. The two estimates are significant at the 10% level in Carvalho et al.'s (2016) analysis.

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Appendix Table 2. Balance Checks

	Mean		
	After Payday	Before Payday	p-value
	(1)	(2)	(3)
Age	56.062	55.836	0.747
Male	0.328	0.340	0.515
Household size	1.935	1.953	0.705
Household head	0.843	0.849	0.706
Children in household	0.162	0.173	0.463
Metropolitan area	0.810	0.799	0.499
Current income	1735.856	1740.043	0.937
Share of payday pay amount relative to current income	0.758	0.765	0.534
Financial strain			
Live from paycheck to paycheck	0.480	0.498	0.388
Caloric crunch	0.473	0.467	0.758
Liquidity constrained	0.500	0.506	0.752
Financial hardship	0.404	0.423	0.332
Marital status			
Married	0.323	0.346	0.213
Divorced	0.277	0.275	0.924
Widowed	0.138	0.140	0.867
Never married	0.263	0.239	0.164
Race			
White	0.756	0.766	0.556
Black	0.110	0.090	0.089*
Hispanic	0.084	0.080	0.736
Other race	0.050	0.064	0.130

Notes: The table continues on the next page.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Appendix Table 2. Continued

	Mean		
	After Payday	Before Payday	p-value
	(1)	(2)	(3)
Employment status			
Working	0.284	0.290	0.744
Unemployed	0.067	0.060	0.521
Disabled	0.191	0.207	0.333
Retired	0.391	0.385	0.771
Other employment status	0.067	0.058	0.330
Education			
Less than high school	0.062	0.064	0.855
High school	0.247	0.260	0.465
Some college	0.419	0.415	0.860
College	0.272	0.261	0.534
Annual household income			
Less than $5,000	0.048	0.048	0.990
Between $5,000 and $10,000	0.094	0.105	0.362
Between $10,000 and $15,000	0.134	0.152	0.193
Between $15,000 and $20,000	0.131	0.109	0.081*
Between $20,000 and $25,000	0.147	0.151	0.802
Between $25,000 and $30,000	0.144	0.143	0.941
Between $30,000 and $35,000	0.143	0.138	0.721
Between $35,000 and $40,000	0.158	0.155	0.787

Notes: $N = 2,480$. The data are from the KnowledgePanel experiment by Carvalho et al. (2016). Columns (1)–(2) show the covariate means for the individuals who are randomly assigned to be surveyed after payday, and before payday, respectively. Column (3) gives the p-values from pairwise t-tests which test whether the difference in means between the before and after payday group for a given covariate is different from zero. The difference in means for the covariate black is also significant at the 10% percent level in Carvalho et al.'s (2016) full sample.

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Appendix Table 3. OLS Average Effect Estimates for the Subgroups Analyzed by Carvalho et al. (2016)

Outcome	Number of correct answers per second	Number of correct answers	Total response time (in seconds)
	(1)	(2)	(3)
Panel A. Subgroup: One payment			
Before payday	0.003	−0.014	−1.500
	(0.008)	(0.663)	(1.332)
Constant	0.419***	41.799***	104.461***
	(0.006)	(0.478)	(0.986)
N	1,265	1,265	1,265
Panel B. Subgroup: Financial Hardship			
Before payday	0.007	0.066	−0.949
	(0.009)	(0.670)	(1.474)
Constant	0.447***	42.638***	99.784***
	(0.007)	(0.490)	(1.050)
N	1,026	1,026	1,026
Panel C. Subgroup: Live paycheck to paycheck			
Before payday	0.012	0.435	−1.737
	(0.008)	(0.602)	(1.284)
Constant	0.441***	42.629***	100.863***
	(0.006)	(0.450)	(0.933)
N	1,213	1,213	1,213
Panel D. Subgroup: Annual household income less than $20,000			
Before payday	0.000	−0.376	−0.321
	(0.009)	(0.756)	(1.522)
Constant	0.424***	41.686***	102.278***
	(0.007)	(0.534)	(1.045)
N	1,020	1,020	1,020
Panel E. Subgroup: Caloric crunch			
Before payday	0.011	0.666	−1.190
	(0.009)	(0.645)	(1.353)
Constant	0.433***	42.040***	101.913***
	(0.006)	(0.482)	(0.982)
N	1,165	1,165	1,165
Panel F. Subgroup: Liquidity constrained			
Before payday	0.013	0.257	−1.753
	(0.008)	(0.619)	(1.388)
Constant	0.437***	42.332***	101.796***
	(0.006)	(0.449)	(0.959)
N	1,248	1,248	1,248

Notes: The data are from the KnowledgePanel experiment by Carvalho et al. (2016). Heteroskedasticity-robust standard errors are in parentheses.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Appendix Table 4. Causal Forest Estimates for Typical Individuals in the Vicinity of the Typical Younger Individual

Age	18	19	20	21	22
Current income	(1)	(2)	(3)	(4)	(5)
Panel A. Outcome: Correct answers per second					
\$400	−0.0520*	−0.0388**	−0.0385**	−0.0434**	−0.0420**
	(0.0272)	(0.0172)	(0.0187)	(0.0198)	(0.0166)
\$425	−0.0545**	−0.0451**	−0.0499***	−0.0418**	−0.0386*
	(0.0224)	(0.0183)	(0.0167)	(0.0198)	(0.0211)
\$450	−0.0545**	−0.0525**	−0.0477***	−0.0415*	−0.0387*
	(0.0226)	(0.0227)	(0.0145)	(0.0225)	(0.0207)
\$475	−0.0620**	−0.0369**	−0.0405**	−0.0380**	−0.0377**
	(0.0252)	(0.0160)	(0.0184)	(0.0188)	(0.0182)
\$500	−0.0396*	−0.0296**	−0.0424*	−0.0396	−0.0350*
	(0.0236)	(0.0140)	(0.0220)	(0.0284)	(0.0212)
Panel B. Outcome: Number of correct answers					
\$400	−2.281	−1.569	−1.572	−1.566	−1.376
	(1.848)	(1.276)	(1.282)	(1.467)	(1.430)
\$425	−2.692*	−2.014	−2.221*	−1.398	−1.263
	(1.488)	(1.697)	(1.233)	(1.581)	(1.103)
\$450	−2.693*	−2.419	−2.253	−1.256	−1.268
	(1.475)	(1.597)	(1.844)	(1.356)	(1.073)
\$475	−3.384*	−1.379	−1.817	−1.143	−1.260
	(2.009)	(1.566)	(1.503)	(0.999)	(1.182)
\$500	−1.241	−0.739	−2.138*	−1.546	−1.453
	(1.273)	(1.445)	(1.252)	(1.814)	(1.282)
Panel C. Outcome: Total response time					
\$400	1.626	1.445	1.424	2.705	2.587*
	(2.758)	(2.925)	(2.883)	(2.179)	(1.466)
\$425	2.286	1.068	1.984	3.255*	2.224
	(1.962)	(2.249)	(1.730)	(1.834)	(1.874)
\$450	2.293	2.062	1.559	2.342	2.232
	(2.022)	(2.327)	(1.515)	(2.366)	(1.894)
\$475	1.442	1.042	0.901	2.480	2.146
	(2.881)	(2.886)	(1.628)	(2.072)	(1.833)
\$500	1.572	1.118	0.759	0.657	1.197
	(1.991)	(2.483)	(2.209)	(1.810)	(1.608)

Notes: The data are from the KnowledgePanel experiment by Carvalho et al. (2016). Standard errors are in parentheses. The table shows causal forest estimates for the effect of the poorer financial circumstances before payday on our three cognition outcomes. For more information, see Appendix E.

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Appendix Table 5. Causal Forest Estimates for Typical Individuals in the Vicinity of the Typical Older Individual

Age	73	74	75	76	77
Current income	(1)	(2)	(3)	(4)	(5)
Panel A. Outcome: Correct answers per second					
\$400	−0.0340	−0.0356*	−0.0361**	−0.0359*	−0.0247
	(0.0226)	(0.0193)	(0.0184)	(0.0199)	(0.0205)
\$425	−0.0359**	−0.0364*	−0.0370**	−0.0367*	−0.0253
	(0.0160)	(0.0199)	(0.0188)	(0.0200)	(0.0186)
\$450	−0.0359**	−0.0365*	−0.0370**	−0.0367*	−0.0253
	(0.0156)	(0.0197)	(0.0187)	(0.0200)	(0.0184)
\$475	−0.0366**	−0.0373**	−0.0378**	−0.0376**	−0.0256
	(0.0142)	(0.0176)	(0.0167)	(0.0183)	(0.0175)
\$500	−0.0369**	−0.0376**	−0.0381**	−0.0379**	−0.0369*
	(0.0156)	(0.0178)	(0.0166)	(0.0182)	(0.0197)
Panel B. Outcome: Number of correct answers					
\$400	−0.850	−1.204	−1.436	−1.529	−0.985
	(1.258)	(1.616)	(1.791)	(1.642)	(1.585)
\$425	−0.840	−1.209	−1.445	−1.543	−1.004
	(1.693)	(1.553)	(1.753)	(1.582)	(1.550)
\$450	−0.831	−1.206	−1.442	−1.540	−0.998
	(1.671)	(1.546)	(1.742)	(1.573)	(1.546)
\$475	−0.786	−1.212	−1.452	−1.551	−0.993
	(1.410)	(1.316)	(1.555)	(1.369)	(1.299)
\$500	−0.778	−1.203	−1.440	−1.537	−1.540
	(1.420)	(1.360)	(1.601)	(1.411)	(1.537)
Panel C. Outcome: Total response time					
\$400	4.971	6.238**	6.195**	6.221**	3.822
	(4.102)	(2.828)	(2.887)	(2.895)	(3.701)
\$425	5.834*	6.375**	6.318**	6.350**	3.932
	(3.266)	(2.804)	(2.874)	(2.954)	(3.893)
\$450	5.856*	6.376**	6.319**	6.351**	3.928
	(3.344)	(2.817)	(2.888)	(2.981)	(3.916)
\$475	6.020*	6.517**	6.461**	6.488**	4.022
	(3.538)	(2.535)	(2.695)	(2.767)	(3.704)
\$500	6.200*	6.701**	6.646**	6.668**	6.372**
	(3.741)	(3.209)	(3.321)	(3.234)	(2.715)

Notes: The data are from the KnowledgePanel experiment by Carvalho et al. (2016). Standard errors are in parentheses. The table shows causal forest estimates for the effect of the poorer financial circumstances before payday on our three cognition outcomes. For more information, see Appendix E.

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.