Social Transfers and Labor Supply: Evidence from South Africa's Old Age Pension*

Kanishka Kacker†

June 24, 2020

Abstract

Does the impact of social transfers on labor supply depend on whether the economy is growing or in recession? I answer this question by analyzing the South African Old Age Pension program, one of the largest social transfer programs in a developing country. Exploiting the age cut-off for pension eligibility in a regression discontinuity design and a major demand shock from a nation-wide recession, I find pension transfers reduce labor supply only in recession-hit sectors. The fall in wages due to the recession results in a substitution effect to favor leisure. This combines with the income effect of receiving pension to depress labor supply. Neither effect alone is, however, significant. Recessions often evoke a policy response to raise social transfers. My finding suggests that transfers at such a time reduce labor supply.

Keywords: Pension, Labor Supply, Panel Data, NIDS
JEL Codes: H23, H55, I38, J22, O15

*I thank Farzana Afridi, Monishankar Bishnu, Sneha Bakshi, Abhiroop Mukhopadhyay, Vikram Dayal, Prabal Ray Chowdhury, Eswaran Somanathan, participants at the 14th Annual Conference on Economic Growth and Development 2018 in Delhi and participants at the 94th Annual Conference of the Western Economic Association 2019 in San Francisco for comments on earlier drafts of this paper. Victor Sulla and Nga Thi Viet Nguyen helped in understanding the National Income Dynamics Survey data.

†Indian Statistical Institute; email: kkacker@isid.ac.in, kanishka.kacker@gmail.com

1 Introduction

How do social transfers affect labor supply? And does the answer depend on whether an economy is growing or is in recession? Classical theory predicts that supply falls with an increase in non-wage income, if leisure is a normal good. In reality, the picture is considerably more complex and uncertain [Baird et al., 2018]. Labor supply mostly changes little in response to transfers [Banerjee et al., 2017] but can change in either direction [Baird et al., 2018]. The question of whether the impact of social transfers depends on the state of the macroeconomy has not received much attention.

To answer this question, I study the South African pension program, one of the largest transfer programs in the world particularly for a developing country. The decision to examine labor supply is motivated by certain features of the South African economy. South Africa has very high levels of unemployment, which hasn't changed much since the end of Apartheid nearly three decades ago: in the last quarter of 2019, it reached 29% [Statistics South Africa, 2020] - the highest in the past decade. Job opportunities are particularly poor for the majority black population [Magruder, 2010], [Kingdon and Knight, 2004]. Identifying if and how large transfer payments impact labor supply in such a scenario is crucial, not just for South Africa alone but for other developing countries contemplating aggressive transfer policies.

Recent work on the South African pension program emphasizes two channels through which pension payments influence labor supply, each in opposing directions. Abel [2019] and Ranchhod [2006] emphasize the income effect, and so labor supply falls with pension receipt. In contrast, Ardington et al. [2009] and Ardington et al. [2016] argue pension incomes stimulate outward labor migration, thus raising labor supply. None of these studies examine whether the impact can vary with the state of the economy.

South Africa suffered a recession that started in 2008 and continued until 2010, its most serious since the end of Apartheid. Using data that covers a period of nine years, from 2008 to 2017, allows me to analyze the effect of pension transfers across recession and economic recovery. Existing studies, such as those cited above, have relied on cross-sections or short panels, which by design cannot uncover such effects. My analysis follows prime-age adults - those between 17 and 59 years old - present in all five survey waves for a consistent time-series comparison. Pensions are available to those over the age of 60: I therefore exclude such workers.

My main result is: in the recession years 2008-2010, weekly hours worked is lower by a statistically significant 28 hours for working age adults co-resident with pension-eligible adults compared to working age adults who do not reside with pension-eligible adults. These effects disappear when the recession fades. I argue these results can be explained through a novel channel in addition to the usual income effect of a transfer: an uncompensated substitution effect triggered by a national recession. Wage reductions in sectors hit hard by a recession raise leisure demand through a substitution effect. Negative substitution effects are well known [Ashenfelter and Heckman, 1974], and can arise from the design of the South African pension program as well [Ranchhod, 2006]. The substitution effect I observe, however, comes from a source external to the program, which clearly has important implications for transfer policy. I conclude, from a series of tests designed to isolate each effect, that the substitution effect interacts with the income effect to lower labor supply; each in isolation exerts a weak influence.

Receiving pension is voluntary for those eligible, and so pension inflow cannot be treated as exogenously assigned to members of households with pensioners. I use the eligibility criterion in the pension program to estimate a regression discontinuity design: adults over the age of 60 are eligible to receive the pension, subject to a means test. I compare hours worked by working age adults residing with household members just above the cut-off age - and therefore eligible for the pension - to working age adults residing with household members just below who are ineligible. Implicitly, I rely heavily on the assumption that pension resources are redistributed within households: this assumption is backed by various studies which find strong evidence of such redistribution in South African households [Duflo, 2003], [Posel et al., 2006], [Ardington et al., 2009].

I also supplement the regression discontinuity design with an instrumental variables specification. In this specification, I use the total household members at or over the pension eligible age as an instrument for residing in a household receiving pension inflow, and find similar effects on hours worked. The main advantage of an instrumental variables specification is it allows me to additionally measure effects over the extensive margin - the probability of being employed - for which I fail to find significant effects.

These results are robust to controls for various confounding factors - age, gender, education, race, household size; as well as different bandwidths used to estimate the treatment effect. Edmonds et al. [2005]; Ardington et al.

[2009] find household structure changes in response to the pension. Ardington et al. [2009] further argue the pension relaxes a childcare constraint stimulating outward migration for work, while Ardington et al. [2016] argue the pension allows for rural men to migrate for work. All of these channels are potential confounders of the estimated relationship between hours worked and pension receipt. Through careful construction of the sample and numerous tests I find these alternate channels are unable to explain the labor supply reduction, and thus do not appear to actually confound the main estimates.

This paper contributes to the literature in four ways. First, it is one of the first to argue that the impact of transfer payments on labor supply can vary depending on the growth of the economy. A common policy response to recessionary conditions is to increase social support transfers. Policy involving such transfers should therefore account for their broader welfare impact during recessionary conditions, if my findings generalize to other settings. Second, in contrast to earlier studies on the South African pension which use cross-sections or short panels, I use a considerably longer span of time to consider the impact of the pension, which enables a comparison from recession to recovery periods. I use panel data on over 37,000 individuals resident in over 13,000 households from five consecutive waves of the National Income Dynamics Survey (NIDS), which is a nationally representative survey of households [Brophy et al., 2018]. Third, this is the first study that formally exploits the discontinuity arising from the pension's eligibility criterion in an explicit regression discontinuity design.[1] A regression discontinuity design

[1] Ranchhod [2006] also uses the discontinuity induced by the pension eligible age, how-

allows for recovering parameters similar to that from a randomized experiment [Lee and Lemieux, 2010]. Fourth, unlike earlier studies, I do not find strong impacts on the probability of employment, but only over the hours worked.

Section 2 describes the structure of the South African pension program. Casual evidence over the impact of the program across recession and recovery periods is shown in section 3. These lead to a simple model of labor supply in and out of a recession, which is sketched out in section 4 and emphasizes the interaction of the income and substitution effect. Section 5 describes the regression discontinuity design. The specific channels through which labor supply may be affected are examined in section 6. The implications of these results are laid out in section 7 which concludes the paper.

2 The South African Old Age Pension and the 2008 Recession

Constructed as a way to support elderly whites who retired from the labor force, the pension program dates from pre-Apartheid days. In the past, the age of eligibility varied by gender: the cut-off age was 60 for women and 65 for men. Between 2008 and 2010, the age eligibility for men fell to 60 [Ralston et al., 2016].

Pension amounts and the maximum level of income for eligible recipients

ever, this appears to take place within a conventional least-squares specification. There is no discussion, for instance, of the choice of bandwidth or polynomial order, or the complications induced by incomplete take-up for which I utilize a fuzzy regression discontinuity design.

of the pension have adjusted upwards over time. In 1993, Case and Deaton [1998] report the maximum benefit was Rand 370 a month. The level of the pension would start adjusting downward when the pre-pension sum of income and the value of assets owned exceeded Rand 90 per month and would go to zero if the sum exceeded Rand 370 a month. At present, the maximum level of income is Rand 6510 per month and of assets owned is Rand 1,115,400 per month. These figures double for married individuals.[2].

The South African Old Age pension is thus a means-tested payment, available at present to anyone over the age of 60 who wishes to apply for it. Hanna and Olken [2018] claim these characteristics - means testing combined with self-selection - yield superior screening. The superior screening is evident in the small amount of leakage in this program: 2.7% of those residing in households with no member eligible for the pension report receiving a pension payment while 78% of those residing in households with at least one member eligible for the pension report receiving a pension. [3] Ranchhod [2006] describes how the means test can incentivize workers to substitute leisure for labor.

The pension payment has increased over time in real terms as well. In December 2016 prices, the median payment in 2008 was 1430 Rand which changes to 1541 Rand in 2017. At the same time, pension take-up (calculated

[2]The source for these numbers is the website maintained by the South African government on the old age pension: https://www.gov.za/services/social-benefits-retirement-and-old-age/old-age-Pension

[3]The latter number is lower than 100% either because some people don't pass the means test, or the costs to getting the payment are too high. The latter possibility doesn't appear to be large. Costs to obtaining or delivering the pension would be higher for rural areas but this does not appear to be a major hurdle: from the 2017 NIDS survey, rural Africans - the poorest racial category and most likely to be located far away from urban centers - report higher rates of pension receipt than urban Africans (30% versus 13%).

as the fraction of those drawing pensions to those eligible to do so) decreases from 90% in 2008 to 78% in 2017. Actual pensions received were quite close to stipulated maximum amounts. In 2017, around 57% of the individuals receiving a pension got Rand 1600 a month, while 32% got Rand 1500 per month: the maximum pension amount was set at Rand 1600 per month, or Rand 1620 per month if older than 75 years during this time.[4]

South Africa went through a recession starting around the middle of 2008 and continuing until at least 2010 [South African Reserve Bank, 2009],[Verick, 2012]. The South African government launched a stimulus package with the aim of boosting demand and jobs: interest rate reductions start in December 2008. Despite this the economy contracted severely in 2009, and it wasn't until the second quarter of 2010 that formal sector employment rose after 6 successive quarters of contraction [South African Reserve Bank, 2010]. Gross domestic product began to build in 2010 led mainly by an increase in public sector hiring; quarter-to-quarter employment by the second quarter of 2010 fell by 2.3% in the private sector [South African Reserve Bank, 2010]. By 2012, gross domestic product grew by 3.5% in the second quarter, and while this could not be sustained, there was no consistent quarter-on-quarter contraction as witnessed during the 2008 recession [South African Reserve Bank, 2017].

Data from NIDS can be used to show how the labor market changes as the economy exits recession. NIDS was launched in 2008, and a new survey is conducted approximately every two years. The first two survey waves were

[4] The source for these numbers is the website maintained by the South African government on social security programs: https://www.ssa.gov/policy/docs/progdesc/ssptw/2016-2017/africa/south-africa.html

carried out in 2008 and 2010 while the next three survey waves run from 2012 to 2017. Comparing the first two survey waves to the next three waves reveals how households adjusted as the economy goes from recession to recovery. In the following section I compare households with pensioners to those without as the economy exits recession. This comparision will help understand how the impact of the pension may vary.

3 Summary Evidence on the Impact of the Pension

The impact of the pension on labor supply can be clearly seen in Figure 4. The graph on the left shows hours worked by adults aged between 17 and 59 years old plotted against the age of the oldest household member, for the first two survey waves (2008 to 2011) when the economy is in recession. Black dots represent the average hours worked in households without a pension-eligible member and grey diamonds represent the average hours worked in households with a pension eligible member. Each dot or diamond represents the average for each distinct age, that is, the width of each bin used to construct the dot or diamond is one year.

We can see a sharp drop in hours worked of approximately 9 hours, just at the point where a co-resident becomes pension eligible. The graph on the right shows the same relationship but for the recovery period (2012 to 2017). Here there is no sharp drop, but a slight increase of a little under 4 hours. The overall pattern to the right of the discontinuity point of the pension eligible age of 60 also differs between these two graphs. During the recession period, average hours worked are considerably lower and fall more steeply than during the recovery period.

Apart from labor supply, demographics and household structure differ significantly across pension and non-pension households. Table 1 shows, using data from NIDS, how this difference changes from recession to recovery periods.

A household where there is at least one member drawing a pension is termed a "pension household". "Non-pension households" are those where there is no member drawing a pension. For each of the two time periods - recession (2008 to 2010) and recovery (2012 to 2017) - the sample average and number of observations are shown. I also show the p-value of a two-tailed test for differences in means between the pension and non-pension household within each time period. The summaries in this table are restricted to adults aged between 17 and 59 years old, so none of the results include data on pensioners themselves.

Looking at labor supply decisions, we see there are large differences between pension and non-pension households irrespective of time period. Labor force participation, hours worked and wages are lower for pension households while rates of discouragement and unemployment are higher. Interestingly, total household income - which includes pension income - is not significantly different between pension and non-pension households during the recession years. Reducing labor supply when pension payments flow into the household appears rational as it does not lower total income, assuming households redistribute resources. Even when the difference becomes statistically significantly different in the post-recession years, total household income is only around 5% lower in pension households. This stands in contrast to all other labor supply variables, for whom the differences are quite large.

Total hours worked appears not very different across pension and non-pension households, which might be surprising given what we have observed in Figure 4. The main difference between what is shown in Table 1 and in Figure 4 is that households in Table 1 are grouped by whether a member receives a pension having applied for it while Figure 4 only shows differences across pension-eligible households. As argued above, application is endogenously determined complicating interpretation. Further, the discontinuity is evident in Figure 4 while it is implicitly smoothed over in Table 1. In the formal econometric model I present below, both issues - the discontinuity and endogenous application - are explicitly addressed.

Comparisons of these labor supply variables across time indicate an emergence from recession. For each household type, labor force participation, wages and household income rise over time. Rates of discouragement and unemployment fall.[5]

Demographic differences arise mostly along the dimensions of race and urbanization. Black South Africans are more likely to draw pensions: reflecting the racial profile of the country most of the survey respondents are black. Rural households are much more likely to have pension recipients. This could be due to rates of job arrival and job quality which are both poor in rural areas; also, rural households might have older household members. Most of these demographic differences change little over the years, apart

[5]However, the narrow type of unemployment does not change much over time.

from education increasing slightly which is perhaps simply due to the sample population maturing.

Household structure responds to the pension payment. Pension households are larger, typically with a larger share of older members and a lower share of working age members. Such a pattern is consistent with younger members choosing to live with older members (or vice versa), or choosing not to leave older members when there is a pensioning member in the household. It will be important to eliminate the influence of household structure when examining the effect of pension transfers on labor supply decisions, as it responds endogenously to pension receipt.

The results in this table are suggestive of very strong differences in labor market outcomes between residents in pensioner households and residents in non-pensioner households. Broadly, labor supply is lower in pensioner households. I sketch below a brief theoretical explanation that explains why this may be the case, and specifically, why the recession may exacerbate these differences. I will then test these theories formally, utilizing the discontinuity induced by the age-cut off for pension eligibility.

4 Causes of the Labor Supply Shifts

A simple model of labor supply can help understand the effects of the pension. A worker chooses between consumption and leisure: more hours spent working raises income thus allowing for increased consumption but lowering leisure. This tradeoff can be illustrated quite simply, using a setup borrowed fromVarian [2014].

Let p denote the price of consumption of a composite good, and C the amount of the composite good. The amount of labor - measured in terms of hours worked - is given by L, and the wage rate is w. With $\bar{L}$ denoting the maximum amount of labor possible, $\bar{L} - L = R$ will be the leisure consumed. The maximum leisure consumed $\bar{R}$ therefore equals $\bar{L}$.

The budget constraint can be written as:

$$pC + wR = p\bar{C} + w\bar{L} \tag{1}$$

The left hand side represents the total value of consumption plus the "value" of leisure - obtained by multiplying the total amount of leisure time by the wage rate, which is the opportunity cost of not working. The right hand side represents the total value of the endowment of consumption and the income from working, or the total of non-work and work income.

The optimal mix of consumption and leisure is shown in Figure 1. Here, consumption is on the vertical axis and leisure is on the horizontal. X represents the original consumption-leisure bundle, the point of tangency between a given set of preferences shown by the indifference curve u_0 and the budget line indicated in the figure. The vertical intercept for this budget line equals $\bar{C} + \frac{w}{p}\bar{R}$.

Pension arrival will increase non-work income, shifting the budget line upward by the increase in non-work income, represented by $\bar{C}' > \bar{C}$. As the pension payment is lump-sum it does not alter the relative return for working, meaning the upward shift will take place parallel to the original budget line. For the given set of preferences, the new point of consumption will be at Y, with an increase in the amount of consumption and an increase in the amount of leisure. The entire effect of the pension on leisure operates via an income effect, and if we assume leisure is a normal good (implicit in the construction of the indifference curves), then demand for leisure will rise. In Figure 1, leisure demand rises from R_0 to R_1, which I denote by A.

The impact of a recession is outlined in Figure 2. The recession hit worker sees a decline in the wage rate, from w to w' with $w' < w$, which reduces the "price" of leisure. From the substitution effect we know this will increase the amount of leisure demanded. This is not the only effect of a change in prices: there is the income effect following from the price change and specific to the case of labor supply, the endowment income effect. Varian [2014] describes the derivation of these effects and shows the end result of these three effects is ambiguous. In terms of the figure the fall in the wage rate flattens the budget line which will pivot the budget line outward, and because a fall in the wage rate means a fall in work income, the pivoted budget set will shift downward. It is unclear what the end effect on leisure demand would be - I have drawn a case representing an increase meaning the substitution effect dominates the income effect thus raising leisure demand but it is straightforward to see how leisure demand may fall. The increase in leisure demand equals B.

Pension arrival during a recession can have effects shown in Figure 3. Here, the budget line pivots and shifts down due to the fall in the wage rate (shown by the dotted line) and the pension shifts the budget line upward. Again it is unclear apriori which effect will dominate to ultimately determine leisure demand. In comparison to the case where only the recession takes place but there is no pension income, it is logical to conclude leisure demand will increase by a greater amount. And if the recession increases the demand for leisure, then again the increase in leisure demand in this case will be greater than the increase just due to the pension alone. For this case, the increase in leisure demand equals C.

For the preferences that would imply indifference curves as shown, leisure demand increases in all three cases: $A > 0, B > 0, C > 0$ with $C > A$ and $C > B$. It is possible that $B < 0$, but $C \geq B$ must always be true if leisure is a normal good. And if leisure is a normal good, $A > 0$. As long as $B \geq 0$, $C \geq A$. That is, the combination of pension arrival and a recession will have larger reductions in labor supply than just a recession or pension arrival alone.

5 Regression Discontinuity Estimates of the Labor Supply Response

The regression discontinuity model I estimate can be written as:

$$Y_{iht} = \mu_0 + \tau T_{iht} + \sum_{j=1}^{p} \mu_{-,j}(X_{iht} - c)^j + \sum_{j=1}^{p} \mu_{+,j} T_{iht}(X_{iht} - c)^j + \boldsymbol{Z'_{iht}\gamma} \tag{2}$$

Here i indexes individual, h indexes household and t indexes the survey wave. Y is the outcome, which is hours worked. X is the running variable, which is the age of the oldest household member, from which we subtract c, the cut-off for deciding pension eligibility which equals 60. T refers to the treatment indicator, i.e. whether the household has a pensioner as a member or not. Finally, Z is a vector capturing a variety of controls.

Equation 2 states the following: we run a weighted least squares regression of the outome on a constant, the treatment indicator, a p-order polynomial on the running variable and the covariates. The weights equal $K((X_{iht} - c)/h)$, where $K()$ is a kernel function, and h the bandwidth.

In order to estimate equation 2, I need to specify a choice of polynomial order p, kernel function K and the bandwidth h. The choice of polynomial order is guided by recent work which suggests higher-order polynomials are likely to be influenced by outlier observations [Gelman and Imbens, 2018]. Throughout I set $p = 1$, implying a linear fit. The choice of a kernel function in practice does not appear to heavily influence estimates, I adopt the triangular kernel for $K()$, which puts greater weight on observations near the cut-off. Bandwidth h is chosen to minimize the mean squared error of the treatment effect and restricted to be symmetric on both sides of the cutoff. Bandwidth choice can heavily influence estimates, so I undertake a robustness check to various alternate choices of bandwidth.

The parameter of interest is τ which captures the change in the outcome at the cut-off. The main identifying assumption is that individuals on either side of the cut-off are essentially the same, differing only in their exposure to treatment. Given the exogenous treatment assignment, treatment can then be taken to be randomly assigned around the cut-off. Therefore the change in outcome at the cut-off can be attributed to the treatment.

As shown above, the pension has both incomplete take-up by eligible individuals (to the right of the cut-off) and limited leakage to non-eligible individuals (to the left of the cut-off). Compliance is therefore imperfect which motivates estimating a fuzzy regression discontinuity model. A fuzzy model splits the estimation into two stages. The first stages estimates the indicator variable for pension take-up as a function of the running variable and the eligibility for a pension, while the second estimates the outcome of interest as a function of the same variables. The ratio between these two gives us our parameter of interest τ [Cattaneo et al., 2018]. Since information on pension take-up is explicitly used in this regression discontinuity design, it addresses both issues of a discontinuous shift and endogenous pension application identified earlier as possible threats to interpretating the pension's effects.

In a regression discontinuity design, continuous variables are easily handled but categorical variables - such as labor force participation or employment indicators - are harder to estimate. For categorical outcomes, I instead estimate an instrumental variables specification described below, which can handle both types of variables easily. To see how robust our main parameter of interest τ is, I will also report the impact on hours worked under both specifications - the regression discontinuity design and the instrumental variables method. Using both specifications will therefore tell us how workers respond on the extensive and intensive margin to pension arrival.

All estimates come from a sample of adults aged between 17 and 59 years. I also impose the constraint that they be present in all five waves of NIDS and in households that don't attrite. These choices are made to address attrition in the NIDS data [Abel, 2019]. All results are therefore subject to the caveat that they are coming off a sample that does not attrite.

Validity of the Regression Discontinuity Design The age an individual declares is central to the validity of this particular discontinuity design. If this variable is manipulated, the main identifying assumption will fail, as the cut-off cannot be treated as exogenously given. Comparing outcomes at the thresholds of the cut-off cannot then be attributed to the treatment alone. There is no reason, however, to suspect that survey respondents would manipulate their age, as they derive no benefit from this.

To understand whether treatment assignment can be taken to be random, I examine a simple frequency plot of the age of the individual. If there is a discrete jump at the age of 60, we can infer some manipulation of the age variable. As Appendix Figure A1 shows, however, there is a smooth and continuous trend in the frequency of the age variable around the cut-off value of 60. More formally, I test for statistical differences between the probability of observing a 59 year old and a 60 year old and fail to reject the null hypothesis of no differences: the p-value for the difference in means is 0.67.

In addition to the running variable, covariates should not change discontinuously at the cut-off for the regression discontinuity design to be valid. If they do, that would mean we are possibly conflating labor supply responses with covariate responses. The first four columns of Appendix Table A1 show the response of two measures of household composition to the pension: the total number of young adults and working age adults. The total number of young adults does rise but only in the first two waves. This is both a comforting and disconcerting result: comforting as other studies establish household structure changing in response to the pension [Edmonds et al., 2005]

[Ardington et al., 2009] and disconcerting because disentangling household composition from pension arrival is now necessary. In the next four columns of Appendix Table A1, I show that neither household size, age nor education levels change discontinuously around the cut-off. Since race and gender are discrete variables, I cannot test for a discontinuous jump in them, which is why they are absent.

Plotting Hours Worked versus Age To first assess whether a discontinuity exists, we should see a sharp jump in hours worked when the oldest household member turns 60. We have already seen Figure 4 that shows this to be the case. Importantly, the drop only appears to occur for the first two waves.

Estimates of the Pension's effect on Labor Supply The first two columns of table 2 show treatment effect estimates of the pension, following the method laid out in equation 2. I include the following individual level controls: gender, race, education, a quadratic in age and household size.

Labor supply falls by a statistically significant 28 hours, but only in the first two waves. The sample mean and standard deviation of hours worked in primary jobs are 38 and 16.5 respectively; therefore, a fall in hours worked by 28 hours is very large. In the following three waves, the effect is much smaller and statistically insignificant. Since I am restricting the sample to be the same set of workers, this response cannot be held accountable to changing types of workers.[6]

[6] Restricting further the sample to be balanced does not alter the results much: while the sample size diminishes considerably the estimate for the first two waves is a statistically

Appendix Table A2 reports first stage estimates for these specifications, confirming that pension receipt is strongly correlated with age. Appendix Table A3 shows results with two additional outcomes: hours worked in all jobs for the salaried alone and hours worked in primary jobs for the salaried.[7]

The robustness of these estimates to bandwidth choice are shown in Appendix Table A4. To examine the role of the cut-off, I have also estimated placebo tests by changing the cut-off to 59: the estimates are statistically insignificant. These results are not shown but are available on request. [8]

To further ascertain robustness of the main estimate, I present the results of a fixed effects instrumental variables model, which addresses endogenous pension take-up by instrumenting for household pension status. The main advantage of using an instrumental variables model is that it allows us to examine decisions over the extensive margin which are discrete by nature: the probability of employment. Discrete outcomes are difficult to analyze using a regression discontinuity.

The instruments I use are similar to those used by Abel [2019], Duflo [2003] or Case and Deaton [1998] - the presence of pension eligible household members, determined by their age. I use the total number of male and female

significant reduction of 20 hours while that for the next three waves is a statistically insignificant 8 hour rise. These results are not shown but are available on request.

[7]Those who are self employed have no hours worked entered under the primary job heading, therefore all these outcomes are outcomes for the salaried.

[8]I have assessed the robustness of the main results to including household level controls - estimates change very little when household level controls are incorporated. As household composition appears to change across the cut-off, I have also tried including measures of household composition - the total number of children (ages 0 to 5), young adults (ages 6 to 17), working age adults (ages 18 to 50) and older adults (ages 51 to 59) as covariates. Results do not change much: τ for the first two waves is a statistically significant 25 hour decrease while for the next three is a statistically insignificant 3 hour decrease - suggesting that household composition changes cannot be the only factor causing supply to fall. These results are not shown but are available on request.

pension eligible household members as instruments for whether a household has a pensioner. That is, I estimate the following specification:

$$Y_{iht} = \beta_1 * Pension\ Household_{iht} + \boldsymbol{Z'_{iht}\gamma} + \lambda_i + \epsilon_{iht} \tag{3}$$

$$\begin{aligned} Pension\ Household_{iht} = \alpha_1 * Total\ Pension\ Eligible\ Males_{ht} \\ + \alpha_2 * Total\ Pension\ Eligible\ Females_{ht} \\ + \boldsymbol{Z'_{iht}\gamma} + \lambda_i + \nu_{iht} \end{aligned} \tag{4}$$

Here, Y_{iht} denotes the labor supply decision of individual i in household h during survey wave t. *Pension Household* is an indicator variable that equals one if i is in a household where a member claims a pension. As before, $\boldsymbol{Z}$ is a vector of controls. Since these are panel data, I employ an individual fixed effects specification, given by λ_i, thus using variation within individuals who witness a change in household pension status between survey waves. This allows for a differencing out of any time invariant unobservable, such as ability, that could possibly confound the estimate of β_1. Standard errors are clustered by a grouped household identifier. [9]

The identifiying assumption here is that pension eligibility (dictated by age which is assigned exogenously) determines the presence of a pensioner, and affects labor supply decisions of working age individuals only through the household pension status conditional on controls $\boldsymbol{Z}$. Household pensioner status is determined by the presence of pension eligible individuals, so we

[9] As individuals can change households across time, their errors are likely to be correlated within each household over time. For this reason, it is also not possible to cluster by any one household identifier since individuals can change households over time. I define, therefore, a grouped household identifier for the string of households generated by each individual's choice of residence in each wave.

expect α_1 and α_2 to be positive.[10]

Columns 3 and 4 of Table 2 show estimates from the instrumental variables specification. Hours worked reduces by a statistically significant 12 hours during the first two waves for those co-resident with a pensioner, as we can see from the coefficient of the household pension indicator variable. In the next three waves, however, this coefficient is much smaller in size and statistically insignificant. First stage results are shown further down. The instruments are strongly correlated with the endogenous household pension variable, which we can tell by the F-statistic and the individual coefficients. Further, the overidentification tests indicate we can reject the instruments being correlated with the second stage error terms.[11]

The main advantage of such a specification is that it allows us to analyze responses over the extensive margin: the probability of being employed. These are shown in Panels A and B of Appendix Table A5. Employment probability is statistically insignificantly related to pension arrival while the instruments appear both valid and relevant.

[10] I distinguish between gender of pensioners: previous work suggests redistribution of pension resources takes place when a female pensioner receives payment [Duflo, 2003], [Posel et al., 2006], [Ardington et al., 2009]. Whether the presence of pension eligible members affects supply decisions only through the household's pension status is perhaps debatable. Pension arrival can shift bargaining powers of elderly women [Ambler, 2016]; we also know that household composition can switch in response to the pension. In either case, measures of household composition can also be affected by pension eligibility, and affect labor supply, making it imperative to include as a control.

[11] Columns 3 and 4 leave out total household members aged between 50 and 59; including them has little effect on the results. I have not reported these results but they are available on request.

6 Channels of the Labor Supply Response

The combination of an income effect and a recession-induced substitution effect induces a labor supply reduction. I undertake multiple tests designed to isolate these channels.

I first present evidence that wages fell for workers in recession hit sectors, which is necessary for a substitution toward leisure to take place. Following Verick [2012], I construct an indicator variable for sectors that are affected by the recession: mining, manufacturing, wholesale and retail trade, and financial, real estate and business services. Then, using the regression discontinuity design I will show hours worked falls for workers in these sectors residing in pensioner households. Thus it is the combined impact of an income effect and a substitution effect that leads to a labor supply reduction.

Table 3 shows how monthly wages change by survey wave and sector in real terms. The time period covered by each survey wave is indicated in parentheses next to the respective wave.[12] I use an individual fixed effects specification, so the specification is identified off of workers switching between sectors over time. The base case is the first survey wave. We see that wages rise over time, in real terms, but workers in recession hit sectors suffered badly due to the recession. In the second wave, workers in sectors not affected

[12] The specification can be stated as follows:

$$w_{i,t} = \beta_0 + \sum_{t=2}^{5} \beta_{t0} * (s_{i,t} \times \lambda_t) + \beta_s * s_{i,t} + \sum_{t=2}^{5} \beta_{t1} * \lambda_t + \Gamma * X_{i,t} + \lambda_i + \epsilon_{i,t}$$

Here w is real wages for individual i at time t, in December 2016 prices. s is an indicator variable for whether i works in a recession hit sector at time t. λ_t is an indicator for which survey wave i is observed in. X is a vector of controls, and Γ the associated coefficient matrix. λ_i is an individual fixed effect.

by the recession have wages that are 5 times those of workers in recession hit sectors. These differences in wages narrow over time as the recession fades. By the fourth wave, the differences in wage rates become statistically insignificant. Although the difference becomes statistically significant again in the fifth wave, the fall is much smaller than that in the second wave. Such a decline in wage rates can reduce labor supply through a substitution effect.

If the reduction in supply comes purely from the substitution effect, we should not see a big change for workers co-resident with pensioners. If, however, only the income effect is operative, we should see large changes for workers co-resident with pensioners, irrespective of whether they are hit by recession or not. If both interact, and only the interaction matters, then we should see large changes only for recession hit workers co-resident with pensioners.

Panel A of Table 4 shows the reduction in hours worked is large and statistically significant during the recession years for workers in sectors hit by the recession who co-reside with pension eligible adults. For workers in non-recession sectors, we cannot reject a null effect at conventional levels of significance.[13] Post-recession, workers in neither sectors show a statistically significant change in labor supply following pension receipt. [14] This pattern of results strongly suggests the combination of an income effect from the

[13]The error is higher for non-recession sector workers despite a larger sample being used to construct the estimate, so it isn't being artifically generated by sample size following the sample selection criterion.

[14]It is possible that some workers choose to leave these sectors in between survey waves during the recession years so that some of the response in the non-recession hit sectors could be coming off such workers. Ideally we would want to focus on workers who do not change sectors or simply look at one survey period: either of these restrictions thins the sample too much to have meaningful estimates.

pension and a substitution effect causes labor supply to change. Each effect by its own does not appear to be large enough to shift labor supply.

In the remaining part of Table 4, I examine two related hypothesis in terms of gender and skill. Since the recession hit sectors are male-dominated we would expect to see larger supply reductions for male workers compared to female workers.[15] I argue below that supply reductions should also be larger for medium skilled workers.

Returns to skill are convex in South Africa: in the fifth survey wave in 2017, wages in high skill jobs were higher by 145% relative to medium skilled jobs. In turn, wages for medium skilled jobs were higher than wages in low skilled jobs by 128%.[16]

Medium skilled workers are more likely to suffer given the nature of the recession in terms of what sectors were affected; low skilled workers would be equally likely to be in a recession hit sector or not. High skill workers are unlikely to suffer drastic reductions in demand. For the overall sample, 65% of workers in recession affected sectors were medium skilled while only 34% of workers in non-affected sectors were medium skilled. Cross-sectoral differences are smaller for high and low skill workers.

Panel B of Table 4 shows results when restricting the sample to either

[15] Workers are 10% more likely to be male in the recession hit sectors for the overall sample.

[16] Skill definitions come from Girdwood and Leibbrandt [2009]. "Low Skill" includes military and elementary occupations. "Medium Skill" includes clerks; service workers, shop, market sales workers; skilled agricultural and fishery workers; craft and related trades workers; and plant, machinery operators and assemblers. "High Skill" includes legislators, senior officials, managers; professionals; technicians and associate professionals. Girdwood and Leibbrandt [2009] define 4 skill levels with technicians and associate professionals given a level in between medium and high. I have included these workers as high skill workers as otherwise the sample would be too thin to run a regression discontinuity for just them.

gender. We see results confirming our expectations: male workers see a large, statistically significant fall during the recession years while for female workers we cannot rule out a null effect.[17] Results from splitting the sample by skill are in Panel C of Table 4. The medium skilled see a large and statistically significant reduction in supply. The estimate for the low skilled is lower and statistically insignificant, while that for the high skilled is implausibly large and statistically insignificant.[18]

Other Mechanisms

Endogenous household composition Household composition changes endogenously in response to the pension, so I present estimates of tests designed to eliminate household composition as a possible explanation of the labor supply response.

One way to deal with endogenous household composition is to eliminate its influence entirely by exploring whether transfers can have effects across households. To do so, I focus on working age individuals who have parents of pensionable age that reside in other households. This proceeds on the assumption that parents who are transferring part or all of their pension to their children will continue to do so even if the children don't live with them. Household composition changes now cannot be traced to a pensioning member of the household. By construction, therefore, endogenous household composition is eliminated as a confounding factor. If we continue to

[17]Again, the sample sizes are comparable between these workers so the larger error cannot be completely driven by the sample size induced by the sample selection condition.

[18]These latter estimates possibly reflect a smaller sample size.

find similar effects of the pension as documented so far, we can safely conclude endogenous household composition cannot explain all of the observed response. [19]

Appendix Table A6 shows estimates when the regression discontinuity is estimated off children who do not reside with their parents. Based on these results, we can rule out household composition as being the sole mechanism through which labor supply adjusts to pension arrival. When there are any pension eligible parents, or pension eligible mothers, labor supply falls in the first two waves. The presence of a pension eligible mother has stronger effects, re-iterating earlier work which finds income from female pensioners having strong effects [Duflo, 2003], [Posel et al., 2006], [Ardington et al., 2009]. All of these effects shrink to statistical insignificance in the next three waves. From these results, we can conclude endogenous household composition cannot explain all the labor supply response.

Out migration and child care constraints Ardington et al. [2009] argue the pension relaxes a child care constraint, allowing working age individuals - particularly mothers - to migrate for work. Ardington et al. [2016] further argue the pension helps fund labor migration for young rural men, particularly those with a matriculate degreee. I test for the influence of these channels in Appendix Table A7. The effects are statistically insignificant when I consider mothers alone (Panel A of Appendix Table A7) or rural men alone (Panel B of Appendix Table A7).[20] While it would be preferable to further split the

[19]In this test, since I do not know whether the non-resident parent takes up a pension or not, I estimate a sharp regression discontinuity design.

[20]Indeed the effect for hours worked by rural men in primary jobs goes in the opposite direction to what we would expect if the pension funds labor migration. The estimated

sample for rural men by educational status, this delivers very small sample sizes. Abel [2019] too considers this mechanism and finds it fails to explain labor supply decisions.

Inter-generational transfers A related aspect of these labor supply responses is whether it takes place across generations, within them or both. I re-estimate the main specification separately for the younger generation - defined to be those aged between 17 and 35 years - and the older generation, defined to be those between 36 and 59 years of age. The younger generation reduces labor supply by a larger amount than the older generation, again only during the first two waves. These results suggest that pension payments amount to an intergenerational transfer when they are redistributed within the household. Appendix Table A8 has these results. Note that in trying to eliminate household composition as an explanatory variable, we implicitly assumed transfers are intergenerational in nature, at least partly. It is therefore comforting to see the same pattern holds within a household as well.

Gendered impacts Earlier work [Duflo, 2003], [Posel et al., 2006], [Ardington et al., 2009] suggests pension incomes are more likely to be allocated to other household members when a female pensioner receives payment. Since a major assumption of the present study is that households redistribute pension resources, it is important to see if this redistribution varies by pensioner gender. I restrict the sample such that there is no female pensioner in the

confidence intervals are however too wide to be meaningful although they are bounded away from zero.

household and compare it to the sub-sample where there is at least one female pensioner in the household. Statistically significant reductions are only observed for households with at least one female pensioner, again for only the first two waves. Appendix Table A9 has these results. Ambler [2016] argues such gender based differences come from a change in relative bargaining powers within the household.

Credit constraints Pension resources can alleviate credit constraints, increasing search in the short run and enabling a labor supply increase at least over the medium-run. The results so far run in the opposite direction, suggesting this channel does not appear operative in South Africa.

Abel [2019] notes that the effect of pensions is unclear for unemployed prime-aged adults. Searching for work can rise following pension inflow but can also lead to an increase in reservation wages. I fail to reject a null effect of the pension on reservation wages. The result for this test is in Appendix Table A10. [21] If search increased, we should observe strong effects on the probability of employment. As the results on the extensive margin (Appendix Table A5) show however, these effects are not strong. Such a pattern is consistent with reservation wages not shifting much in response to pension arrival.

If credit constraints matter, Blattman et al. [2014] suggests cash transfers can enable new businesses to start up. This implies labor supply should

[21] Here I report two outcomes - reservation wages and fair wages. Reservation wages are responses by survey respondents to the question "What is the absolute lowest take-home wage that you would accept for any permanent, full-time work (per-month)?". Fair wages are responses to the question "What do you think would be a fair take-home monthly wage for you, given your age, education and skills?"

respond positively to pension inflow for the self-employed. Examining the impact of pensions solely for the self-employed reveals a statistically insignificant effect, whether examining those co-resident with pensioners or resident in households separate from the pension eligible parents (Appendix Table A11). Credit constraints do not appear to be binding even for self-employed individuals in South Africa.

7 Conclusion

In this paper, I demonstrate pensions affect labor supply of prime aged adults co-resident with pensioners through a combination of an income and a substitution effect. The pension raises non-wage income, which will cause a reduction in labor supply, assuming leisure to be a normal good. A recession hits some sectors of the economy more than others, and reduces wages. This triggers a substitution effect that can increase leisure demand. Hours worked reduces by 28 hours - nearly 1.7 standard deviations - in response to pensions, only for workers in recession-hit sectors at the time when the economy was in a recession. Neither channel - the income or substitution effect - is powerful enough to bring down labor supply by itself. Importantly, the probability of being employed appears unaffected by pension arrival. I implicitly assume households redistribute pension resources: this is backed by earlier studies [Duflo, 2003], [Posel et al., 2006], [Ardington et al., 2009] as well as additional analysis I have carried out and described in the section above.

There at least two implications of these results. First, labor supply is unlikely to reduce in response to pension payments when demand is healthy. The implication is large unconditional cash transfers may not always have distortionary effects. Identifying when these distortions are likely to arise can be a way to ensure that a cash transfer program is effectively implemented.

Second, it is precisely when the economy is in a recession that demands to raise transfers will increase. If the results documented here hold more generally, policies aimed at transfer programs should then adjust for the expected fall in labor supply. It is unclear whether the labor supply reduction is welfare reducing or welfare improving, but what is clear is that welfare impacts fall on a larger group of people during recessionary events. Given the vast economic crisis likely to occur on a global scale due to Covid-19, this implication perhaps bears great importance.

A major point of difference between this paper and other studies of the South African pension program is that I find employment probabilities to be unaffected; only the intensive margin of labor supply responds to pension arrival. This is possibly because the labor market in South Africa is very tight, and leaving a job might just be too risky.

Ultimately these pension payments are transfers from South Africa's rich to their poor. All transfers have to be funded by taxation, which inevitably entails a deadweight loss. A full calculation of these effects remain to be carried out and is beyond the scope of the paper.

References

Abel, M. (2019). Unintend labor supply effects of cash transfer programmes: Evidence from South Africa's Old Age Pension. *Journal of African Economies*, ejz009.

Ambler, K. (2016). Bargaining with grandma: The impact of the South African Pension on household decision-making. *Journal of Human Resources 51*(4), 900–932.

Ardington, C., T. Bärnighausen, A. Case, and A. Menendez (2016). Social protection and labor market outcomes of youth in South Africa. *ILR Review 69*(2), 455–470.

Ardington, C., A. Case, and V. Hosegood (2009). Labor supply responses to large social transfers: Longitudinal evidence from South Africa. *American economic journal: Applied economics 1*(1), 22–48.

Ashenfelter, O. and J. Heckman (1974). The Estimation of Income and Substitution Effects in a Model of Family Labor Supply. *Econometrica 42*(1), 73–85.

Baird, S., D. McKenzie, and B. Ozler (2018). The effects of cash transfers on adult labor market outcomes. *IZA Journal of Development and Migration 8*(1), 22.

Banerjee, A. V., R. Hanna, G. E. Kreindler, and B. A. Olker (2017). Debunking the stereotype of the lazy welfare recipient: Evidence from cash transfer programs. *The World Bank Research Observer 32*(2), 155–184.

Blattman, C., N. Fiala, and S. Martinez (2014). Generating Skilled Self-Employment in Developing Countries: Experimental Evidence from Uganda. *Quarterly Journal of Economics 129*(2), 697–752.

Brophy, T., N. Branson, R. C. Daniels, M. Leibbrandt, C. Mlatsheni, and I. Woolard (2018, 1). *National Income Dynamics Study panel user manual* (1 ed.). Cape Town: Southern Africa Labour and Development Research Unit.

Case, A. and A. Deaton (1998). Large cash transfers to the elderly in south africa. *The Economic Journal 108*(450), 1330–1361.

Cattaneo, M. D., N. Idrobo, and R. Titiunik (2018). A practical introduction to regression discontinuity designs: Volume I. Monograph prepared for Cambridge Elements: Quantitative and Computational Methods for Social Science. Cambridge University Press.

Duflo, E. (2003). Grandmothers and Granddaughters: Old-Age Pensions and

Intrahousehold Allocation in South Africa. *The World Bank Economic Review 17*(1), 1–25.

Edmonds, E. V., K. Mammen, and D. L. Miller (2005). Rearranging the family? Income support and elderly living arrangements in a low-income country. *Journal of Human Resources 40*(1), 186–207.

Gelman, A. and G. Imbens (2018). Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business and Economic Statistics*, 1–10.

Girdwood, S. and M. Leibbrandt (2009). Intergenerational mobility: Analysis of the NIDS Wave 1 Dataset. Discussion Paper No. 15.

Hanna, R. and B. A. Olken (2018). Universal Basic Incomes versus Targeted Transfers: Anti-Poverty Programs in Developing Countries. *Journal of Economic Perspectives 32*(4), 201–226.

Kingdon, G. G. and J. Knight (2004). Unemployment in South Africa: The nature of the beast. *World Development 32*(3), 391–408.

Lee, D. S. and T. Lemieux (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature 48*(2), 281–355.

Magruder, J. R. (2010). Intergenerational networks, unemployment, and persistent inequality in south africa. *American Economic Journal: Applied Economics 2*(1), 62–85.

Posel, D., J. A. Fairburn, and F. Lund (2006). Labour migration and households: A reconsideration of the effects of the social pension on labour supply in South Africa. *Economic Modelling 23*(5), 836–853.

Ralston, M., E. Schatz, J. Menken, F. X. Gomez-Olive, and S. Tollman (2016). Who Benefits—Or Does not—From South Africa's Old Age Pension? Evidence from Characteristics of Rural Pensioners and Non-Pensioners. *International Journal of Environmental Research and Public Health 13*(85), 1–14.

Ranchhod, V. (2006). The effect of the South African Old Age Pension on labour supply of the elderly. *South African Journal of Economics 74*(4), 725–744.

South African Reserve Bank (2009). Quarterly Bulletin. No. 253.

South African Reserve Bank (2010). Quarterly Bulletin. No. 258.

South African Reserve Bank (2017). Quarterly Bulletin. No. 286.

Statistics South Africa (2020). Quarterly Labor Force Survey.

Varian, H. R. (2014). *Intermediate Microeconomics: A Modern Approach* (9 ed.). WW Norton and Company.

Verick, S. (2012). Giving up Job Search during a Recession: The Impact of the Global Financial Crisis on the South African Labour Market. *Journal of African Economies 21*(3), 373–408.

Figures

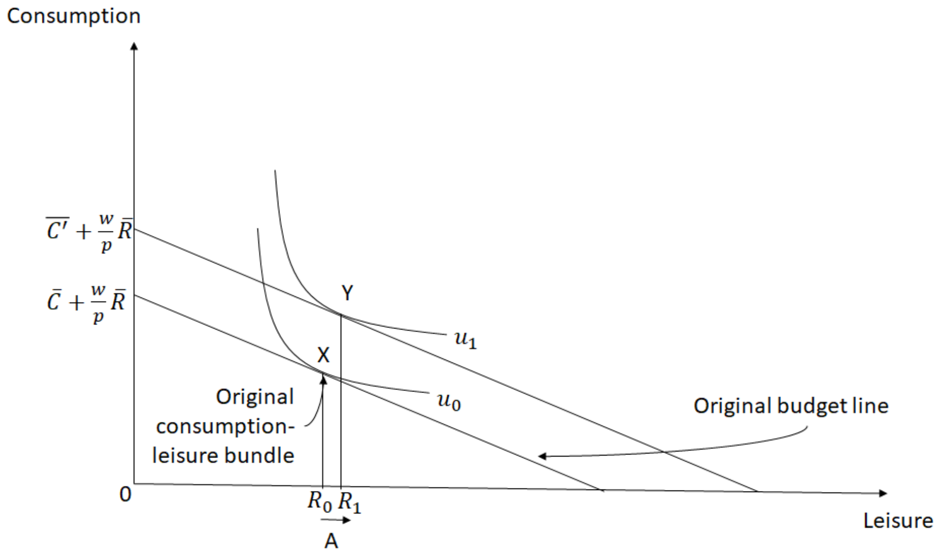

Figure 1: Pension Arrival Alone

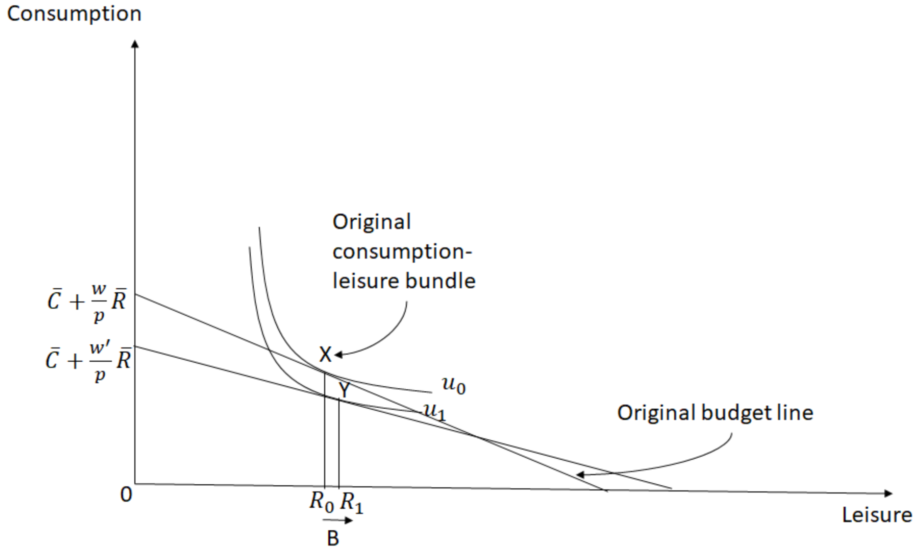

Figure 2: Recession Alone

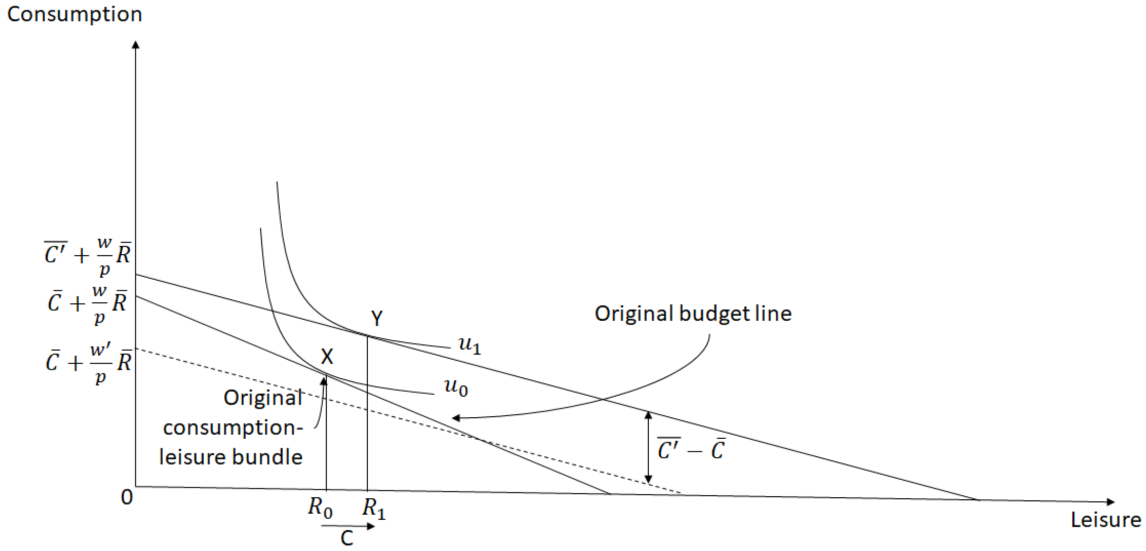

Figure 3: Pension Arrival × Recession

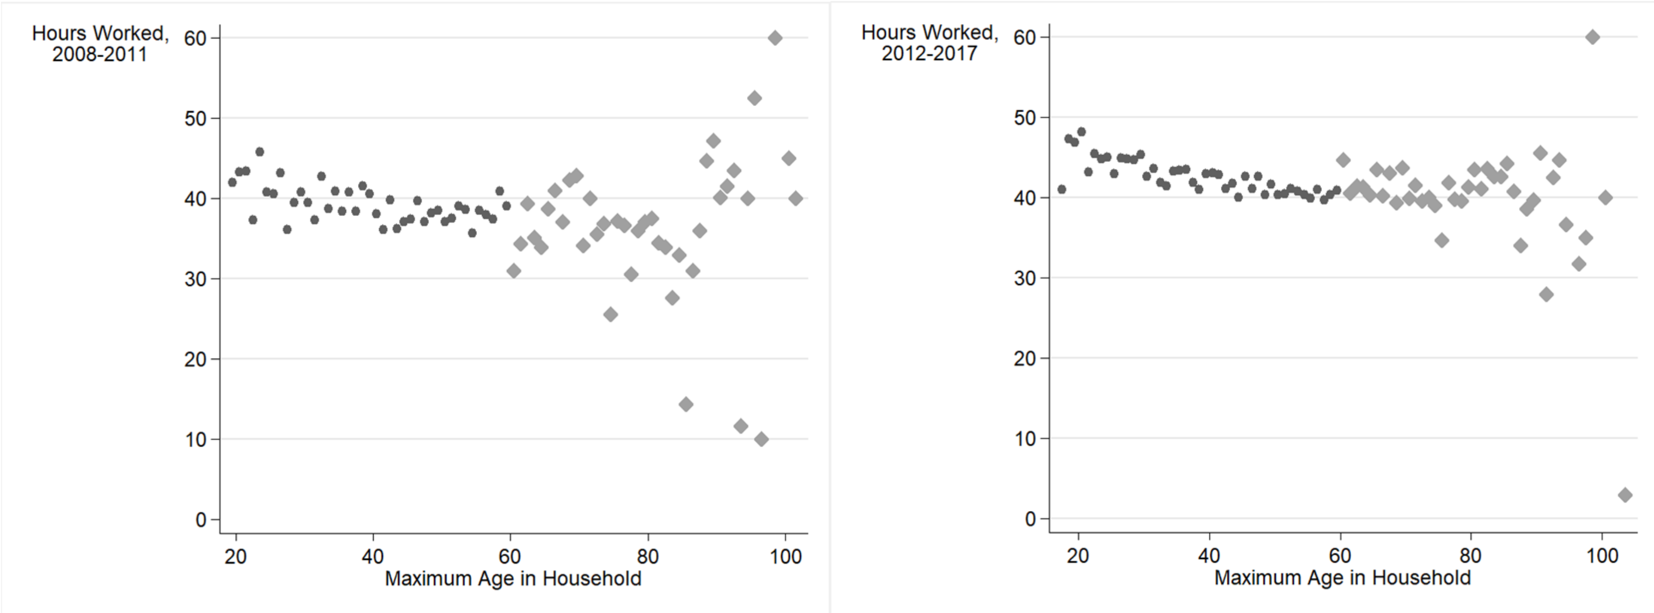

Figure 4: This figure shows hours worked by adults between 17 and 59 years of age, against the age of the oldest person resident in the household. Black dots represent average hours worked in households without a pension-eligible household member. Grey diamonds represent average hours worked in households with a pension-eligible household member. Each dot or diamond represents the average for a distinct age: the width of each bin is 1 year. Hours drop by about 9 hours just at the point when a co-resident becomes pension eligible, but only in the first two waves. In the next three waves, no drop occurs; indeed, the hours worked increase by a smaller amount. (Source: National Income Dynamics Survey Waves 1 to 5)

Table 1: Summary Statistics: Labor Supply, Demographics and Household Structure

	Waves 1 and 2: 2008 - 2011					Waves 3 to 5: 2012 - 2017				
	Non Pension Household	# Obs	Pension Household	# Obs	p	Non Pension Household	# Obs	Pension Household	# Obs	p
Panel A: Labor Supply										
Labor Force Participant*	.552	13,217	.424	4,310	0	.621	24,938	.477	8,189	0
Discouraged Worker*	.106	8,164	.157	2,165	0	.038	16,097	.07	4,202	0
Broad Unemployment*	.317	8,164	.467	2,165	0	.263	16,097	.42	4,202	0
Narrow Unemployment*	.236	7,300	.369	1,826	0	.234	15,480	.376	3,909	0
Hours Worked †	38.4	3,850	37.6	679	.246	42	8,356	40.9	1,537	.003
Wages ‡	3,769	6,327	2,031	1,676	0	4,552	11,792	2,655	2,818	0
Household Income ‡	6,832	13,821	6,707	4,542	.534	8,175	23,044	7,719	7,247	.009
Panel B: Demographics										
Black	.811	15,148	.852	5,044	0	.832	25,394	.846	8,414	.004
Male	.425	15,148	.419	5,044	.448	.443	25,394	.442	8,413	.847
Urban	.508	15,107	.348	5,042	0	.534	25,394	.363	8,414	0
Years of Education	8.73	15,087	8.63	5,027	.102	9.6	25,313	9.36	8,383	0
Panel C: Household Structure										
Household Size	5.21	15,148	7.28	5,044	0	4.83	25,390	7.26	8,414	0
Fraction 0-5 years	.12	15,148	.119	5,044	.646	.109	25,390	.118	8,414	0
Fraction 6-17 years	.242	15,148	.242	5,044	.766	.218	25,390	.234	8,414	0
Fraction 18-50 years	.532	15,148	.396	5,044	0	.566	25,390	.399	8,414	0
Fraction 51+	.106	15,148	.243	5,044	0	.107	25,390	.248	8,414	0

Sample consists of working age adults (17 to 59 years old), who are present in all survey rounds, and reside in households that do not attrite.

*: Denote those not active in the labor force by n, unemployed who have stopped looking for work by d, unemployed but looking for work by u and those employed by e. Labor Force Participant $= (u+e)/(n+d+u+e)$; Discouraged Worker $= d/(d+u+e)$; Broad Unemployment $= (d+u)/(d+u+e)$ and Narrow Unemployment $= u/(u+e)$.

†: Hours worked in the primary job are reported; ‡: Wages and household income are in real terms, with December 2016 as the base year.

Table 2: Labor Supply and Pension Arrival

Outcome: Hours Worked in a Week at the Primary Job				
	(1)	(2)	(3)	(4)
Estimation Technique	Regression Discontinuity		Instrumental Variables	
Period	2008-2011	2012-2017	2008-2011	2012-2017
Parameter				
Discontinuity Estimate (τ)†	-28.42	2.40		
	(8.89)	(3.62)		
Household Pension Indicator			-12.57	-0.35
			(4.51)	(1.63)
Outcome Mean	38.26	41.84	38.26	41.84
Outcome Standard Deviation	16.51	13.69	16.51	13.69
Controls	Y^a	Y^a	Y^b	Y^b
Observations	4,502	9,868	2,178	7,021
Effective Observations:				
Left of Cut-off	1556	1936	-	-
Right of Cut-off	512	874	-	-
First Stage Estimates				
Total Pension Eligible Males	-	-	0.080	0.435
			(0.06)	(0.040)
Total Pension Eligible Females	-	-	0.658	0.615
			(0.07)	(0.033)
First Stage F-Statistic			47.00	307.90
Overidentification Test Statistic			0.10	3.32
p-value Overidentification			0.76	0.07

Estimates come from a sample restricted to working age adults aged between 17 and 59. Standard errors are in parentheses, clustered by household for the regression discontinuity estimate and by a grouped household identifier for the instrumental variables estimate. See the text for details.
† Point estimates and standard errors incorporate an estimated bias term in calculating the treatment effect. For details of the regression discontinuity design, refer to the text.
(a) Controls include gender and race of the individual; urbanization status of the household the individual resides in; years of education for the individual; a quadratic in the age of the individual and household size.
(b) Controls include an individual fixed effect; household size; age and square of age; total household residents aged 0 to 5, 6 to 17, and 18 to 50; years of education and square of years of education; an indicator for urbanization status; household head's age and square of age; household head's years of education; and an indicator for whether the household head is female.

Table 3: Wage Rates and the 2008-2010 Recession

Wave #2 (2010/11)	3,128.187
	(1127.962)
Wave #3 (2012)	2,818.617
	(1193.578)
Wave #4 (2014/15)	4,227.595
	(1971.132)
Wave #5 (2017)	5,359.840
	(2547.185)
Recession Hit Sector	1,207.544
	(533.648)
Wave #2 × Recession Hit Sector	-2,542.074
	(906.781)
Wave #3 × Recession Hit Sector	-1,147.981
	(492.309)
Wave #4 × Recession Hit Sector	-649.72
	(619.221)
Wave #5 × Recession Hit Sector	-1,214.318
	(587.402)
Controls[a]	Y
Individual Fixed Effects	Y
Observations	13,345
R-squared	0.015
# Individuals	5,749
Outcome Mean	3312
Outcome Standard Deviation	7330

The outcome here is monthly wages, in December 2016 prices. Estimates come from a sample restricted to working age adults aged between 17 and 59. Standard errors are in parentheses, clustered by a grouped household identifier.
(a) Controls include an individual fixed effect; household size; age and square of age; total household residents aged 0 to 5, 6 to 17, and 18 to 50; years of education and square of years of education; an indicator for urbanization status; household head's age and square of age; household head's years of education; and an indicator for whether the household head is female.

Table 4: Interpreting the Impact of the Pension Program on Labor Supply

Panel A: Sectors Hit by Recession

	(1)	(2)	(3)	(4)
Sample	Recession Sector		Non-Recession Sector	
Period	2008-2011	2012-2017	2008-2011	2012-2017
Discontinuity	-20.80	-3.697	-39.17	8.388
Estimate	(9.76)	(5.79)	(24.68)	(7.28)
Controls	Y	Y	Y	Y
Observations	1,645	3,826	2,671	5,868

Panel B: Gender of Recipient

	(1)	(2)	(3)	(4)
Sample	Male		Female	
Period	2008-2011	2012-2017	2008-2011	2012-2017
Discontinuity	-32.87	-1.449	-39.06	6.478
Estimate	(13.46)	(6.61)	(21.43)	(6.196)
Controls	Y	Y	Y	Y
Observations	2,188	4,774	2,314	5,094

Panel C: Skill Levels

	(1)	(2)	(3)	(4)	(5)	(6)
	Low Skill		Medium Skill		High Skill	
Period	2008-2011	2012-2017	2008-2011	2012-2017	2008-2011	2012-2017
Discontinuity	-17.49	7.055	-27.62	1.616	-191.4	4.706
Estimate	(27.28)	(6.93)	(12.47)	(5.64)	(114.7)	(12.74)
Controls	Y	Y	Y	Y	Y	Y
Observations	1,461	3,356	2,076	4,310	824	1,859

Estimates come from a sample restricted to working age adults aged between 17 and 59. Standard errors are in parentheses, clustered by household. Discontinuity estimates and standard errors incorporate an estimated bias term in calculating the treatment effect.

Controls for Panels A through C, and regression discontinuity specifications, are the same as in columns 1 and 2 in Table 2.

Figure A1: This figure shows the total count of NIDS responders by age, with the height of each column indicating the total number of responders for each age. Those under the age of 60 are shown in black colored columns while the grey colored columns show those at least the age of 60. The absence of any bunching at the age of 60 indicates age is unlikely to have been manipulated. The assignment of age as the running variable in a regression discontinuity design is therefore valid.

Table A1: Examining Changes in Covariates

	Household Composition							
	Young Adults, Ages 6 to 16		Working Age Adults, Ages 17 to 59		Household Size		Education	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Period	2008-2011	2012-2017	2008-2011	2012-2017	2008-2011	2012-2017	2008-2011	2012-2017
Point Estimate†	0.70	-0.20	-0.39	-0.25	0.11	0.18	-0.19	-0.40
Standard Error†	0.34	0.17	0.70	0.22	1.49	0.60	0.75	0.38
Lower 95% CI†	0.03	-0.54	-1.77	-0.68	-2.80	-0.99	-1.65	-1.14
Upper 95% CI†	1.36	0.14	0.99	0.18	3.03	1.36	1.28	0.34
Outcome Mean	1.58	1.45	2.61	2.52	5.73	5.44	8.70	9.54
Outcome Standard Deviation	1.51	1.55	1.56	1.62	3.26	3.44	3.76	3.44
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Observations	20,071	33,692	20,071	33,692	20,071	33,692	20,071	33,692
Effective Observations:								
Left of cut-off	5193	7559	2229	4773	5193	6667	4635	5700
Right of cut-off	2712	4434	1464	3262	2712	4063	2490	3676

Estimates come from a sample restricted to working age adults aged between 17 and 59, from a fuzzy regression discontinuity design with a household declaring a pension recipient as the outcome in the first stage. Standard errors are clustered by the household the individual resides in. A local linear polynomial with a triangular kernel is used to construct the estimates in the neighborhood of the cut-off; the bandwidth is selected by minimizing the mean squared error of the treatment effect. The same bandwidth is chosen on both sides of the cut-off. Controls for columns (1) to (4) include gender and race of the individual; urbanization status of the household the individual resides in; years of education for the individual; a quadratic in the age of the individual and household size. In columns (5) and (6) household size is dropped as a control. Education is dropped as a control for columns (7) and (8).
† Robust point estimates, standard errors and confidence intervals incorporate an estimated bias term in calculating the treatment effect.

Table A2: First Stage Estimates: Fuzzy Regression Discontinuity Design

Second Stage Outcome		Hours Worked in a Week (Primary Job)			Hours Worked in a Week (All Jobs, Salaried Only)			Hours Worked in a Week (Primary Job, Salaried Only)	
First Stage Outcome		Household has Pensioner			Household has Pensioner			Household has Pensioner	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Waves:	2008-2017	2008-2011	2012-2017	2008-2017	2008-2011	2012-2017	2008-2017	2008-2011	2012-2017
Point Estimate†	0.35	0.27	0.42	0.36	0.24	0.41	0.36	0.27	0.42
Standard Error†	0.04	0.06	0.04	0.04	0.07	0.04	0.04	0.06	0.04
Lower 95% CI†	0.28	0.15	0.33	0.28	0.10	0.32	0.28	0.15	0.34
Upper 95% CI†	0.43	0.39	0.50	0.43	0.38	0.50	0.43	0.39	0.50
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	14,370	4,502	9,868	14,161	4,424	9,737	14,151	4,415	9,736
Effective Observations									
Left of cut-off	2837	1556	1936	2800	1155	1912	2796	1524	2176
Right of cut-off	1237	512	874	1223	432	864	1222	504	927

Estimates come from a sample restricted to working age adults aged between 17 and 59, from a fuzzy regression discontinuity design with a household declaring a pension recipient as the outcome in the first stage. Standard errors are clustered by the household the individual resides in. A local linear polynomial with a triangular kernel is used to construct the estimates in the neighborhood of the cut-off; the bandwidth is selected by minimizing the mean squared error of the treatment effect. The same bandwidth is chosen on both sides of the cut-off. Controls include gender and race of the individual; urbanization status of the household the individual resides in; years of education for the individual; a quadratic in the age of the individual and household size.

† Robust point estimates, standard errors and confidence intervals incorporate an estimated bias term in calculating the treatment effect.

Table A3: The Pension Program and Labor Supply: Regression Discontinuity Estimates

	Hours Worked in a Week (Primary Job)			Hours Worked in a Week (All Jobs, Salaried Only)			Hours Worked in a Week (Primary Job, Salaried Only)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Period	2008-2017	2008-2011	2012-2017	2008-2017	2008-2011	2012-2017	2008-2017	2008-2011	2012-2017
Point Estimate†	-4.84	-28.42	2.40	-5.74	-40.93	2.59	-4.24	-28.58	2.85
Standard Error†	4.09	8.89	3.62	4.15	13.45	3.89	3.99	9.13	3.54
Lower 95% CI†	-12.85	-45.84	-4.70	-13.88	-67.30	-5.03	-12.06	-46.47	-4.09
Upper 95% CI†	3.16	-11.00	9.50	2.40	-14.57	10.21	3.58	-10.68	9.79
Outcome Mean	40.71	38.26	41.84	41.29	39.19	42.24	40.71	38.22	41.84
Outcome Standard Deviation	14.73	16.51	13.69	15.75	18.55	14.20	14.65	16.46	13.59
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	14,370	4,502	9,868	14,161	4,424	9,737	14,151	4,415	9,736
Effective Observations:									
Left of cut-off	2837	1556	1936	2800	1155	1912	2796	1524	2176
Right of cut-off	1237	512	874	1223	432	864	1222	504	927

Estimates come from a sample restricted to working age adults aged between 17 and 59, from a fuzzy regression discontinuity design with a household declaring a pension recipient as the outcome in the first stage. Standard errors are clustered by the household the individual resides in. A local linear polynomial with a triangular kernel is used to construct the estimates in the neighborhood of the cut-off; the bandwidth is selected by minimizing the mean squared error of the treatment effect. The same bandwidth is chosen on both sides of the cut-off. Controls include gender and race of the individual; urbanization status of the household the individual resides in; years of education for the individual; a quadratic in the age of the individual and household size.

† Robust point estimates, standard errors and confidence intervals incorporate an estimated bias term in calculating the treatment effect.

Table A4: Sensitivity of Main Estimates to Bandwidth Choice

Panel A: First Two Waves (2008-2011)	Bandwidth chosen to minimize: Mean Squared Error of Regression Discontinuity Estimate				
	Common around cutoff[a]	Different around cutoff[b]	Common, Sum of Estimates[c]	Minimum of Difference and Sum[d]	Median of Different, Common, Common Sum[e]
Point Estimate†	-28.42	-37.08	-29.7	-29.7	-29.77
Lower 95% CI†	-45.84	-60.2	-48.09	-48.09	-47.93
Upper 95% CI†	-11	-13.97	-11.32	-11.32	-11.62
Standard Error†	8.886	11.79	9.381	9.381	9.264
Controls	Y	Y	Y	Y	Y
Observations	4502	4502	4502	4502	4502
Effective Observations:					
Left of cut-off	1556	1717	1556	1556	1556
Right of cut-off	512	437	512	512	512
Outcome Mean			38.26		
Outcome Standard Deviation			16.51		

Sensitivity of Main Estimates to Bandwidth Choice: Table A4 Continued

Panel B: First Two Waves (2008 - 2011)	Bandwidth chosen to minimize: Coverage Error of Confidence Intervals				
	Common around cutoff[a]	Different around cutoff[b]	Common, Sum of Estimates[c]	Minimum of Difference and Sum[d]	Median of Different, Common, Common Sum[e]
Point Estimate†	-34.77	-44.91	-35.99	-35.99	-35.79
Lower 95% CI†	-61.49	-83.5	-64.15	-64.15	-63.54
Upper 95% CI†	-8.056	-6.323	-7.822	-7.822	-8.036
Standard Error†	13.63	19.69	14.37	14.37	14.16
Controls	Y	Y	Y	Y	Y
Observations	4,502	4,502	4,502	4,502	4,502
Effective Observations:					
Left of cut-off	1037	1037	901	901	1037
Right of cut-off	404	279	363	363	363
Outcome Mean			38.26		
Outcome Standard Deviation			16.51		

Sensitivity of Main Estimates to Bandwidth Choice: Table A4 Continued

Panel C: Last Three Waves (2012 - 2017)	Bandwidth chosen to minimize: Mean Squared Error of Regression Discontinuity Estimate				
	Common around cutoff[a]	Different around cutoff[b]	Common, Sum of Estimates[c]	Minimum of Difference and Sum[d]	Median of Different, Common, Common Sum[e]
Point Estimate†	2.399	2.08	2.401	2.399	2.399
Lower 95% CI†	-4.699	-4.116	-4.613	-4.699	-4.618
Upper 95% CI†	9.498	8.276	9.415	9.498	9.416
Standard Error†	3.622	3.161	3.578	3.622	3.58
Controls	Y	Y	Y	Y	Y
Observations	9,868	9,868	9,868	9,868	9,868
Effective Observations:					
Left of cut-off	1936	2523	1936	1936	1936
Right of cut-off	874	1000	874	874	874
Outcome Mean			41.84		
Outcome Standard Deviation			13.69		

Sensitivity of Main Estimates to Bandwidth Choice: Table A4 Continued

Panel D: Last Three Waves (2012 - 2017)	Bandwidth chosen to minimize: Coverage Error of Confidence Intervals				
	Common around cutoff[a]	Different around cutoff[b]	Common, Sum of Estimates[c]	Minimum of Difference and Sum[d]	Median of Different, Common, Common Sum[e]
Point Estimate†	3.479	2.778	3.376	3.479	3.375
Lower 95% CI†	-6.638	-5.72	-6.6	-6.638	-6.602
Upper 95% CI†	13.6	11.28	13.35	13.6	13.35
Standard Error†	5.162	4.336	5.09	5.162	5.091
Controls	Y	Y	Y	Y	Y
Observations	9,868	9,868	9,868	9,868	9,868
Effective Observations					
Left of cut-off	1161	1421	1161	1161	1161
Right of cut-off	620	713	620	620	620
Outcome Mean			41.84		
Outcome Standard Deviation			13.69		

The outcome for all estimates is weekly hours worked in the primary job. Estimates come from a sample restricted to working age adults aged between 17 and 59, from a fuzzy regression discontinuity design with a household declaring a pension recipient as the outcome in the first stage. Standard errors are clustered by the household the individual resides in. A local linear polynomial with a triangular kernel is used to construct the estimates in the neighborhood of the cut-off. Robust point estimates, standard errors and confidence intervals incorporate an estimated bias term in calculating the treatment effect. Controls for all the estimates include gender and race of the individual; years of education for the individual; a quadratic in the age of the individual and household size. Notes: (a) Bandwidth constructed to be the same on either side of the cutoff (oldest household member is at least 60 years old and thus pension eligible); (b) Bandwidth constructed to be different on either side of the cutoff; (c) Bandwidth constructed to be the same on either side of the cutoff, but the estimator whose mean squared error (Panel A) or whose coverage error (Panel B) is being minimized is the sum of the regression coefficients on either side of the cutoff not the difference as in (a) and (b) above; (d) The lower bandwidth value comparing between bandwidth values calculated in (a) and (c); (e) This is the bandwidth which takes the median value amongst the bandwidths calculated in (a), (b) and (c)

†: Robust point estimates, standard errors and confidence intervals incorporate an estimated bias term in calculating the treatment effect.

Table A5: Extensive Margin Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:	Probability (Employed)					
Period	2008-2017		2008-2011		2012-2017	
Panel A: OLS						
Pension Household	-0.032	-0.025	-0.054	-0.044	-0.031	-0.026
	(0.008)	(0.008)	(0.018)	(0.019)	(0.011)	(0.012)
Observations	48,376	47,118	17,143	16,464	31,233	30,654
R-squared	0.098	0.1	0.014	0.017	0.064	0.067
# Individuals	12,327	12,288	9,241	9,199	11,859	11,829
Panel B: IV						
Pension Household	-0.030	-0.025	-0.057	-0.032	-0.022	-0.022
	(0.014)	(0.016)	(0.040)	(0.048)	(0.019)	(0.021)
Observations	47,083	45,772	15,804	14,530	29,861	29,200
# Individuals	11,034	10,942	7,902	7,265	10,487	10,375
Individual Controls †	Y	Y	Y	Y	Y	Y
Household Controls ‡	N	Y	N	Y	N	Y
Individual Fixed Effects	Y	Y	Y	Y	Y	Y

Extensive Margin Estimates: Table A5 Continued

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:	Probability (Employed)					
Period	2008-2017		2008-2011		2012-2017	
First Stage Instruments						
Total Pension Eligible Males	0.350	0.335	0.269	0.245	0.366	0.364
	(0.015)	(0.016)	(0.033)	(0.034)	(0.023)	(0.023)
Total Pension Eligible Females	0.624	0.572	0.590	0.535	0.604	0.571
	(0.013)	(0.012)	(0.028)	(0.030)	(0.021)	(0.017)
First Stage F-Statistic	1826	1878	343	243.5	680.3	1023
Overidentification Test Statistic	3.787	0.031	0.326	0.237	2.75	0.321
p-value Overidentification	0.052	0.86	0.568	0.626	0.097	0.571

Estimates come from a sample restricted to working age adults between the ages of 17 and 59, who are present in all five NIDS waves and reside in non-attrition households. Standard errors clustered by grouped household identification in parentheses. As individuals can move between households, I construct a group identification code that takes on a unique value for the string of household identification numbers formed by combining all five wave houeshold identification numbers.

†: Individual level controls include household size; age and square of age; total household residents aged 0 to 5, 6 to 17, and 18 to 50; years of education and square of years of education; and an indicator for urbanization status.

‡: Household level controls include household head's age and square of age; household head's years of education; and an indicator for whether the household head is female.

IV estimates come from a GMM model. The first stage F-Statistic is the Kleibergen-Paap Wald F statistic, calculated to account for clustered standard errors. Hansen's J statistic is used to calculate the overidentification test statistic.

Table A6: Eliminating Household Composition as a Confounding Factor

	(1)	(2)	(3)	(4)
Sample	Only Non-Resident Pension Eligible		Only Female Non-Resident Pension Eligible[a]	
Period	2008-2011	2012-2017	2008-2011	2012-2017
Discontinuity	-8.65	0.32	-25.67	-3.81
Estimate	(2.96)	(1.26)	(8.43)	(3.36)
Controls[b]	Y	Y	Y	Y
Observations	4,502	9,868	904	2,430

Estimates come from a sample restricted to working age adults aged between 17 and 59. Standard errors are in parentheses, clustered by household. Discontinuity estimates and standard errors incorporate an estimated bias term in calculating the treatment effect.
(a): Outcomes for columns (1) and (2) is hours worked in the reported primary job, and for columns (3) and (4) the outcome is hours worked in all jobs. Self-employed individuals are included in the latter outcome.
(b): Controls, and regression discontinuity specifications, are the same as in the first two columns of Table 2.

Table A7: Other Possible Mechanisms for the Labor Supply Effect

Panel A: Childcare Constraint - Mothers Only	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:	Hours Worked, Primary Job			Hours Worked, All Jobs		
Period	2008-2017	2008-2011	2012-2017	2008-2011	2008-2011	2012-2017
Point Estimate†	0.55	-799.30	10.04	19.13	-37.10	30.42
Standard Error†	9.59	447.70	6.39	22.68	34.22	21.97
Lower 95% CI†	-18.24	-1677.00	-2.49	-25.32	-104.20	-12.65
Upper 95% CI†	19.35	78.26	22.57	63.58	29.97	73.48
Outcome Mean	38.07	35.89	39.13	38.28	35.97	39.57
Outcome Standard Deviation	14.83	16.89	13.60	26.94	32.57	23.13
Controls	Y	Y	Y	Y	Y	Y
Observations	5,630	1,828	3,802	7,514	2,685	4,829
Effective Observations						
Left of cut-off	1371	485	1059	1371	1005	885
Right of cut-off	769	253	568	892	455	597

Other Possible Mechanisms for the Labor Supply Effect: Table A7 Continued

Panel B: Labor Migration - Rural Men						
	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:	Hours Worked, Primary Job			Hours Worked, All Jobs		
Period	2008-2017	2008-2011	2012-2017	2008-2017	2008-2011	2012-2017
Point Estimate†	-10.00	-60.18	2.06	-10.52	-57.99	3.49
Standard Error†	9.99	23.55	6.58	9.52	32.04	11.51
Lower 95% CI†	-29.58	-106.30	-10.84	-29.18	-120.80	-19.08
Upper 95% CI†	9.58	-14.01	14.96	8.15	4.81	26.06
Outcome Mean	43.58	40.66	44.96	43.56	42.36	44.17
Outcome Standard Deviation	14.91	17.38	13.37	26.78	33.27	22.71
Controls	Y	Y	Y	Y	Y	Y
Observations	2,659	856	1,803	3,735	1,269	2,466
Effective Observations						
Left of cut-off	585	323	452	1014	411	520
Right of cut-off	267	95	205	429	148	270

Estimates come from a sample restricted to working age adults aged between 17 and 59, from a fuzzy regression discontinuity design with a household declaring a pension recipient as the outcome in the first stage. Standard errors are clustered by the household the individual resides in. A local linear polynomial with a triangular kernel is used to construct the estimates in the neighborhood of the cut-off; the bandwidth is selected by minimizing the mean squared error of the treatment effect. The same bandwidth is chosen on both sides of the cut-off. Controls include gender and race of the individual; urbanization status of the household the individual resides in; years of education for the individual; a quadratic in the age of the individual and household size.

† Robust point estimates, standard errors and confidence intervals incorporate an estimated bias term in calculating the treatment effect.

Table A8: Intergenerational Transfers Within Households

	(1)	(2)	(3)	(4)	(5)	(6)
Period	2008-2017		2008-2011		2012-2017	
	Ages 17 to 35	Ages 36 to 59	Ages 17 to 35	Ages 36 to 59	Ages 17 to 35	Ages 36 to 59
Robust Point Estimate	-7.13	2.14	-38.93	-28.10	0.53	8.74
Robust Standard Error	5.01	6.38	16.26	22.10	4.45	9.16
Robust Lower 95% CI	-16.95	-10.36	-70.81	-71.42	-8.18	-9.21
Robust Upper 95% CI	2.68	14.63	-7.06	15.22	9.25	26.69
Outcome Mean	41.86	39.61	39.26	37.44	42.91	40.72
Outcome Standard Deviation	14.82	14.56	17.10	15.96	13.65	13.65
Controls	Y	Y	Y	Y	Y	Y
Observations	7,038	7,332	2,018	2,484	5,020	4,848
Effective Observations Left	988	1585	409	494	668	730
Effective Observations Right	770	416	270	121	548	225

The outcome here is hours worked in the primary job. Estimates come from a sample restricted to working age adults aged between 17 and 59, from a fuzzy regression discontinuity design with a household declaring a pension recipient as the outcome in the first stage. Standard errors are clustered by the household the individual resides in. A local linear polynomial with a triangular kernel is used to construct the estimates in the neighborhood of the cut-off; the bandwidth is selected by minimizing the mean squared error of the treatment effect. The same bandwidth is chosen on both sides of the cut-off. Controls include gender and race of the individual; urbanization status of the household the individual resides in; years of education for the individual; a quadratic in the age of the individual and household size.

† Robust point estimates, standard errors and confidence intervals incorporate an estimated bias term in calculating the treatment effect.

Table A9: Gendered Impacts of the Pension on Labor Supply

	(1)	(2)	(3)	(4)
	At least One Female Pensioner		At least One Male Pensioner	
Period	2008-2011	2012-2017	2008-2011	2012-2017
Point Estimate†	-35.46	3.47	-112.80	-0.08
Standard Error†	12.67	5.36	67.64	8.63
Lower 95% CI†	-60.30	-7.03	-245.40	-17.00
Upper 95% CI†	-10.62	13.98	19.75	16.83
Outcome Mean	38.33	41.84	39.05	42.10
Outcome Standard Deviation	16.48	13.68	29.46	22.10
Controls	Y	Y	Y	Y
Observations	4,390	9,560	5,414	10,864
Effective Observations				
Left of cut-off	1306	1671	1842	2832
Right of cut-off	412	637	270	548

Estimates come from a sample restricted to working age adults aged between 17 and 59, from a fuzzy regression discontinuity design with a household declaring a pension recipient as the outcome in the first stage. Standard errors are clustered by the household the individual resides in. A local linear polynomial with a triangular kernel is used to construct the estimates in the neighborhood of the cut-off; the bandwidth is selected by minimizing the mean squared error of the treatment effect. The same bandwidth is chosen on both sides of the cut-off. Controls include gender and race of the individual; urbanization status of the household the individual resides in; years of education for the individual; a quadratic in the age of the individual and household size.

† Robust point estimates, standard errors and confidence intervals incorporate an estimated bias term in calculating the treatment effect.

Table A10: Individual Preferences: Reservation and Fair Wages

	(1)	(2)	(3)	(4)
Outcome	Reservation Wage		Fair Wage	
Period	2008-2011	2012-2017	2008-2011	2012-2017
Point Estimate†	905	-3364	77.67	295
Standard Error†	2142	2770	3999	1030
Lower 95% CI†	-3294	-8793	-7760	-1724
Upper 95% CI†	5104	2064	7916	2314
Outcome Mean	3506	5252	4366	7797
Outcome Standard Deviation	10198	18379	15807	13020
Controls	Y	Y	Y	Y
Observations	4,854	12,739	5,014	28,144
Effective Observations				
Left of cut-off	853	2237	1175	5373
Right of cut-off	492	1071	723	3178

Estimates come from a sample restricted to working age adults aged between 17 and 59, from a fuzzy regression discontinuity design with a household declaring a pension recipient as the outcome in the first stage. Standard errors are clustered by the household the individual resides in. A local linear polynomial with a triangular kernel is used to construct the estimates in the neighborhood of the cut-off; the bandwidth is selected by minimizing the mean squared error of the treatment effect. The same bandwidth is chosen on both sides of the cut-off. Controls include gender and race of the individual; urbanization status of the household the individual resides in; years of education for the individual; a quadratic in the age of the individual and household size.

†: Robust point estimates, standard errors and confidence intervals incorporate an estimated bias term in calculating the treatment effect.

Table A11: Effect of the Pension on the Self Employed

	(1)	(2)	(3)	(4)
	Co-Resident		Non-Resident	
Period	2008-2011	2012-2017	2008-2011	2012-2017
Point Estimate†	8.045	-0.723	31.06	11.18
Standard Error†	32.67	17.63	38.60	22.21
Lower 95% CI	-55.99	-35.27	-44.59	-32.35
Upper 95% CI	72.08	33.82	106.7	54.71
Outcome Mean	40.98	41.78	40.98	41.78
Outcome Standard Deviation	67.64	41.22	67.64	41.22
Observations	1,868	2,759	1,868	2,759
Effective Observations				
Left of cut-off	387	891	400	758
Right of cut-off	179	367	199	402

Estimates come from a sample restricted to working age adults aged between 17 and 59 who declare themselves to be self-employed. For the first two columns, a fuzzy regression discontinuity design with a household declaring a pension recipient as the outcome in the first stage is used. For the next two columns, a sharp discontinuity design is used as the pension status of the non-resident pensioner is not known. The outcome is hours worked in a week.

Standard errors are clustered by the household the individual resides in. A local linear polynomial with a triangular kernel is used to construct the estimates in the neighborhood of the cut-off; the bandwidth is selected by minimizing the mean squared error of the treatment effect. The same bandwidth is chosen on both sides of the cut-off. Controls include gender and race of the individual; urbanization status of the household the individual resides in; years of education for the individual; a quadratic in the age of the individual and household size.

†: Robust point estimates, standard errors and confidence intervals incorporate an estimated bias term in calculating the treatment effect.