Social transfers and labor supply: Long run rvidence from South Africa

Kacker, Kanishka

Indian Statistical Institute, Delhi

November 2019

Social Transfers and Labor Supply: Long Run Evidence from South Africa

Kanishka Kacker*

PRELIMINARY AND INCOMPLETE. DO NOT CITE.

Abstract

How do large social transfers affect labor supply? This study analyses the South African pension program to answer this question. I exploit a major demand shock - the South African recession that began in 2008 - in a regression discontinuity design to find prime aged adult labor supply falls in response to pension arrival in the household only during the recession for sectors and types of workers affected by the recession. Post-recession, these workers witness an increase in demand and respond by increasing supply. Pension payments consequently have small and statistically insignificant effects on labor supply, a result that contrasts starkly with all existing studies. I argue these results stem from the combination of two forces. When labor demand is weak, the opportunity cost of leisure falls and workers demand more leisure. If a household member draws a pension, with leisure being a normal good, leisure demand increases further.

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

I thank Sneha Bakshi, Farzana Afridi, Prabal Ray Chowdhury, Serena Trucchi, Eswaran Somanathan, Swati Virmani, Abhiroop Mukhopadhyay, Vikram Dayal, 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, Delhi. email: kkacker@isid.ac.in, kanishka.kacker@gmail.com

1 Introduction

How do social transfers affect labor supply? This paper examines the South African pension program to answer this question. The South African pension program is one of the largest transfer programs in existence, particularly for a developing country. The size and scale of these transfers sets it apart from social transfer programs in most developing countries [Case and Deaton, 1998]: in 2017, a third of the monthly household income in South Africa came from the pension program, according to the National Income Dynamics Survey. Social transfer programs as a way to reduce poverty in developing countries are now under increased attention [Hanna and Olken, 2018]. Given redistribution from one household member to another, a policy such as the pension program can have unwanted impacts as transfers meant for the old are distributed to younger household members.

I find pension payment leads to a reduction in labor supply but only during the recession that began in 2008. Further, this fall is limited only to workers and sectors affected by the recession. Labor supply is unaffected by pension payments into the household in the years post-recession. All these effects are for prime-aged adults co-resident with pensioners; pensioners themselves are excluded. Therefore, all changes in labor supply as a result of the pension come from a reallocation of pension income. I argue these responses come from two channels. When labor demand weakens, the opportunity cost of leisure falls and so leisure demand rises. With leisure being a normal good, the arrival of a pension income raises the demand for leisure further. When these two effects operate together, labor supply falls dramatically; individually, neither effect exerts a strong influence.

As the sample is restricted to prime aged adults observed in all waves, the impact of the pension therefore varies over time for the same set of workers. That the pension has a time-varying impact is new as is the interpretation that I propose. Importantly, pension payments can have negligible effects on labor supply. Many studies find large impacts of this program on labor supply in either direction and propose various mechanisms to explain their results [Bertrand et al., 2003] [Ranchhod, 2006] [Ardington et al., 2009] [Ardington et al., 2016] [Abel, 2019]. None of these explanations can rationalize a null effect; my explanation can. The implication is that large social transfers may not leak onto unwanted recipients.

Using the South African recession that began in 2008 as a shock to labor demand helps uncover this mechanism: when the recession is active, weekly hours worked is lower by a statistically significant 28 hours in pensioner households compared to non-pensioner households. These effects disappear when the recession fades.[1] Confirming the demand-led hypothesis I find strong supply reductions only in the sectors which were hit particularly hard by the recession - mining, manufacturing, wholesale and retail trade, and financial services [Verick, 2012]. As most of these sectors are dominated by men, I also find statistically significant and large reductions for men. The supply reduction is concentrated amongst the medium-skilled: I argue below these are the workers most likely to suffer during a recession. Post-recession, workers are able to transition from low to medium skill jobs indicating a

[1] Given the standard deviation in hours worked is 17 hours during the recession period and 14 in the period after, these estimates of the labor supply drop during the recession are very large.

strengthening of demand. Labor force participation also rises in parts of the country where barriers to employment are quite large, offering further evidence of strengthening demand.

To estimate the impact of pension arrival on labor supply, I use a regression discontinuity design which arises from the design of the pension program. Adults over the age of 60 are eligible to receive the pension (subject to a means test) but must apply for it. As the decision to receive the pension is left to the pension eligible, pension inflow cannot be treated as exogenously assigned to members of households with pensioners. I thus utilize the age cut-off of 60 to estimate a regression discontinuity, comparing working age adults residing with just eligible household members to working age adults residing with household members just below the cut-off age of 60.[2] The identifying assumption is that individuals on either side of the cut-off are essentially the same, differing only in their exposure to "treatment" i.e. having a pension earning older household member. Given treatment assignment is exogenous as it derives from the age of an individual, treatment can then be taken to be randomly assigned around the cut-off. Therefore the change in hours worked at the cut-off can be attributed to the treatment.

The main distinguishing feature of this study is the use of long-run panel data. I use information from five rounds of the National Income Dynamics Survey (NIDS) which spans nine years from 2008 to 2017. It is meant to be nationally representative, so we can understand the average impact of the pension program across the country. NIDS follows people over time,

[2] An instrumental variables specification, using exogenous variation induced by the age at which pension eligibility occurs, reveals hours drop by a statistically significant 12 hours between 2008 and 2011-12 and a statistically insignificant small increase from 2013 onward.

while gathering information on employment and pension receipt amongst many other variables so we can see how labor supply responds to pension receipt for the same individual over a nine year period. [3] To the best of my knowledge, this is the first time such data have been used to examine the labor supply response to the pension program. The advantage of long run panel data lies in its ability to reveal how transfer impacts vary over time for the same worker, which is obviously impossible with the short panels [Ardington et al., 2009], [Abel, 2019] or cross-sections [Bertrand et al., 2003]; [Ranchhod, 2006] that have been used so far.[4]

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. Household structure has been demonstrated to change in response to the pension [Edmonds et al., 2005]; [Ardington et al., 2009]; I am able to reject this as a confounding factor. I am also able to reject alternate explanations of labor supply responses to pension receipt: that it relaxes a childcare constraint [Ardington et al., 2009] or that it allows for rural men to migrate for work [Ardington et al., 2016].

The decision to look at labor supply is motivated by certain features of the South African economy. South Africa has high levels of unemployment, which hasn't changed much since Apartheid's end nearly three decades ago. As is well known, South Africa has striking racial disparity in income levels, the nature of the labor market and the distribution of unemployment. In

[3] The first survey was held in 2008 and a new round is held every 2 years. In the first wave, a little over 7,000 households comprising around 26,000 individuals were interviewed. In subsequent waves, these numbers grow.

[4] [Ardington et al., 2016] use a long panel but it is restricted to one district which implies we cannot generalize from their findings to the whole country.

2017, white South Africans earned a little over 8 times what Black South Africans earned, in terms of nominal per-capita household income. Black South Africans were 5% less likely to be employed and 13% more likely to be in informal occupations. Unusually for a developing country, the informal sector is relatively small [Kingdon and Knight, 2004]: 3/4ths of jobs were in the formal sector in 2017 - making a 13% difference large.[5] Even when Black South Africans are able to work, therefore, this is concentrated in a small sector of the economy marked by greater uncertainty over period of work, worker protection and benefits. Identifying the effect of large transfer payments in such a scenario is crucial.

As an empirical matter, labor supply decisions are easily observed at the individual level. Consequently, we can examine labor supply choices as a function of changes in household resources in a transparent way. Other consumption decisions are typically observed only at the household level, making it difficult to understand within household redistribution. The next section briefly describes the pension program and the recession that hit in 2008. I then present some summary evidence of the impact of the pension. A regression discontinuity design - my main specification - then follows. I supplement these estimates with additional tests and conclude with a discussion of these results.

[5] The definition of formality I use is if the agreement to work was a written contract.

2 The South African Old Age Pension: a brief

The pension program dates back from pre-Apartheid days, wherein it was constructed as a way to support elderly Whites who retired from the labor force. It is a means-tested payment, available at present to anyone over the age of 60. In the past, the age of eligibility varied by gender: the cut-off age was 60 for women and 65 for men. In the period from 2008 to 2010, the age eligibility for men fell incrementally to 60 [Ralston et al., 2016]. The decision to receive the pension is up to the pension eligible: the South African pension program can thus be understood as a means-test system together with self-selection, which [Hanna and Olken, 2018] claim yields superior screening.[6]

Pension amounts and the maximum level of income for eligible recipients 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.[7].

[6]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. 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%).

[7]The source for these numbers is the website maintained by the South African government on the old age pension: https://www.gov.za/services/

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 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 these stipulated 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.[8]

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]. In response to this, 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 re-

social-benefits-retirement-and-old-age/old-age-Pension

[8]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

cession [South African Reserve Bank, 2017].

Evidence on the drastic effects of the recession on the labor market can be seen in the NIDS data. In terms of survey waves, the recession would involve the first two waves: the first was held in 2008 and the second between 2010 and 2011. The next three waves from 2012 to 2017 would be the post recession waves. Comparing sample means from these first two survey waves to the next three waves , broad unemployment falls from 33% to 29%, rates of discouragement fall from 11% to 4% and participation in the labor force increases from 53% to 59% for adults aged between 17 and 59 years.[9] Both wages and household incomes rise in real terms, again comparing sample means in the first two survey waves to the next three waves. Wages increase from 3200 Rand to 4034 Rand, an increase of 25% while incomes rise from 7190 Rand to 8465 Rand, an increase of nearly 18%.[10]

3 Summary Evidence on the Impact of the Pension

A program as generous as the pension is likely to involve significant changes on many aspects of individual and household decision making. Table 1 shows how labor supply, demographics and household structure change during and post-recession, for pension and non-pension households. A household where there is at least one member drawing a pension is termed a "pension house-

[9] 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 $= (d+u)/(d+u+e)$.

[10] All figures are in December 2016 prices.

hold". "Non-pension households" are those where there is no member drawing a pension. For each of the two time periods, the sample average and number of observations are shown. I also include 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: the exception to this is total hours worked, which we will examine in more detail.

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.[11]

[11] The narrow type of unemployment does not change much however over time.

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

Finally, we see evidence that households respond to the pension payment by arranging themselves to take advantage of it. 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. It is not clear if these differences hold once we control for potential confounders. Most importantly the decision to apply for and receive the pension is endogenous to the pensioner, and thus the pensioner household. In the following section, I show results using a regression discontinuity design: such a specification allows for a clean comparison between pensioner and non-pensioner households. Before that, however, I sketch a brief theoretical explanation that sets up the main hypothesis.

4 Causes of the Labor Supply Shifts

A simple model of an optimal leisure-labor tradeoff can help understand the effects of the pension on labor supply. A worker can be understood as choosing 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, and much of the following discussion borrows from Varian [2014] in its setup.

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 (implict 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: there is the income effect 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, however in comparison to just the recession case 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.

In the way I have drawn the budget sets and indifference curves, 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. 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. 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. 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.

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.

As argued 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 therfore 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].

Hours worked is the outcome Y: 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 am therefore eliminating the influence of pension receipt on pensioners: any change in labor supply as a result of pension arrival is therefore due to a reallocation of the pension income. 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 data. Attrition is an issue with this survey, taking place systematically [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 clearly 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 (in the first two waves alone, again). 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. I will demonstrate that houeshold composition changes do not explain all of the labor supply responses in a later section. 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, I construct a regression discontinuity plot of hours worked by an individual against the age of the oldest member of the household the individual resides in. We should see a sharp jump when the oldest household member turns 60. Figure 4 shows this to be the case, with hours worked dropping when a household member becomes pension eligible. Importantly, this only appears to occur for the first two waves.

In the left panel of Figure 4, I show a scatter plot of binned means of hours worked by the age of the oldest household member for the first two waves, together with a fitted regression specification drawn across the bandwidth used to estimate the regression discontinuity. Hours at the cut-off age of 60 - when an individual becomes pension eligible - drop when a household gets a pensionable resident. In the panel on the right, which is plotted for the the next three waves, there is no drop. Instead, hours worked appear to rise slightly at the cut-off, while the fitted regression lines imply a much more modest effect.

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. For most of the paper, all regression results come off a sample of working age adults aged between 17 and 59 who are present in all four NIDS waves and reside in households that are not lost through attrition.

Labor supply falls by a statistically significant 28 hours, but only in the first two waves. To understand the size of this estimate, the sample mean and standard deviation of hours worked in primary jobs are 38 and 16.5 respectively. In the following three, 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.[12]

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.[13]

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. [14]

To ascertain how robust the main estimate is, I present the results of a fixed effects instrumental variables model, which addresses endogenous pension take-up by instrumenting for household pension status. 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. Specifically, I use the total number of male and female pension eligible household members as instruments for whether a household has a

[12]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 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.

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

[14]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.

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. [15]

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

[15] 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.[16]

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 which vanishes in the next three. 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 correlared with the second stage error terms.[17]

The main advantage of such a specification is that we can see how responses over the extensive margin - the probability of being employed - is affected. These are shown in Panels A and B of Appendix Table A5. The labor supply response on this margin is statistically insignificantly related to pension arrival while the instruments appear both valid and relevant.

6 Mechanism of the Labor Supply Response

In this section I try to draw out the mechanism underlying the response of labor supply to pension payments. Household composition changes endoge-

[16] 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.

[17] 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.

nously 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. I then consider why the response weakens over time. The main hypothesis is that labor supply reduces when labor markets are tight, that is, during the recession. Workers have difficulty in finding a job and if there is a pension resource to draw upon, work part-time jobs relying on the pension income to make up the short fall in total income. Once the recession is past, and finding jobs is easier, they switch back to full time work and do not need the pension income support. I undertake multiple tests designed to isolate this mechanism.

Endogenous Household Composition Panel A of Table 3 shows estimates constructed from sub-samples selected to filter out endogenous household composition as a confounding factor. By focusing on working age individuals who have parents of pensionable age that reside in other households, I am able to explore whether transfers affect labor supply decisions across households - not just within, which has been the focus so far. 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.[18]

Based on the results in Panel A, we can rule out household composition as being the sole mechanism through which labor supply adjusts to pension arrival. Whether there are any pension eligible parents, or pension eligible mothers, labor supply falls in the first two waves. The presence of a pension

[18] 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.

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.[19] From these results, we can conclude endogenous household composition cannot explain all the labor supply response.

Recession To examine whether the recession induces the reduction in labor supply, I exploit heterogeneity in the labor market over three dimensions: sector, gender and skill. First, 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. If slackening demand for labor causes workers to work part-time jobs relying on the pension income to make up for the shortfall in total household income, then the labor supply response we observe should be strong for sectors affected by the recession. That is, if demand side factors cause supply to fall, then we should observe this fall only in sectors where plausibly labor demand fell.

[19]The outcome variable for the results in columns 3 and 4 of Panel C in Table 3 is weekly hours worked in all jobs: the difference between this outcome and the one I use for much of the paper - hours worked in primary jobs - is that the self employed are included in the former outcome. Gender of the non-resident parent matters more, therefore, for the self-employed and the self-employed are more likely targets for redistribution. Labor supply reductions are weaker for hours worked in primary jobs when there is any non-resident eligible parent than for hours worked in all jobs when there is a female non-resident eligible mother. There is a statistically insignificant effect on hours worked in all jobs when there is any non-resident eligible parent and for hours worked in primary jobs when there is a non-resident eligible mother. Importantly, the overall pattern of a muted response in waves 3 to 5 and a stronger response in waves 1 and 2 remains irrespective of outcome or parent gender. A full description of results including those for hours worked in primary jobs is given in Panel A of Appendix Table A6.

Panel B of Table 3 shows that indeed this is what happens: the reduction in hours worked is large and statistically significant during the recession years for workers in sectors hit by the recession. For workers in non-recession sectors, we cannot reject a null effect at conventional levels of significance.[20] Post-recession, workers in neither sectors show a statistically significant change in labor supply following pension receipt. [21]

Second, since these sectors are male-dominated we would also expect to see larger supply reductions for male workers compared to female workers. [22] Panel C of Table 3 shows results when restricting the sample to either 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.[23]

Third, I consider skill levels of 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%.[24]

---

[20]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.

[21]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.

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

[23]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.

[24]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

High skill workers are unlikely to suffer drastic reductions in demand during a recession. Medium skilled are more likely to suffer given the nature of the recession in terms of what sectors were affected; low skilled workers would be - at least, relatively - equally likely to be in a recession hit sector or not. 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. Similar cross-sectoral differences are smaller for the other two skill levels. Results from splitting the sample by skill are in Panel D of Table 3.

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.[25]

Thus far I have presented evidence consistent with the fall in labor supply as a consequence of pension arrival arising from recession induced demand shocks. Workers affected by the demand shock cut their supply - that is, working part-time jobs - if they have a pension income to rely on. From Table 1, we know for these workers that their household income falls only by a small amount. For workers who cannot draw upon a pension income in their household, and work increased hours as a result, the return cannot therefore be much higher per-hour of work than those workers who are able to rely on a pension transfer. If they could rely on a pension, workers who

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.

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

faced a tight demand market chose part-time work, presumably as these jobs are less desirable but easy to find. If jobs were desired, workers would not wish to only contribute part-time.[26]

One implication of this demand-led hypothesis is that workers, particularly in recession hit sectors, ought to see a rise in skill level over time. Recall we are observing workers initially in the throes of a recession that I contend are later able to transition into better jobs. Increases in skill are correlated strongly with wage increases, and would be broadly consistent with the idea of workers choosing more desirable jobs.

The results from testing this implication are shown in Column 1 of Panel D in Table 3. The outcome here is defined as equal to 1 if the job is of medium skill and 0 if of low skill. To test for how skill levels change over time, I define a wave indicator variable that equals 1 if the observation is in the fourth or fifth survey wave and 0 if in the first three. I use a fixed effect logit model to estimate the relationship between this outcome and an interaction term which consists of the indicator for the recession hit sectors and the wave indicator variable. The sign on the interaction term reveals that skill levels increase over time for those in the recession hit sectors.[27]

In the other columns in Panel D, I also report how responses over the extensive margin vary over time for "Bantustans". These are provinces which

[26]One way to interpret "less desirable" is in terms of wages earned. Using the regression discontinuity and instrumental variable estimates to bound the true impact of hours worked, I define an indicator variable that equals 1 if hours worked is between 12 and 28 hours and 0 otherwise. Regressing wages on this variable with a host of controls reveals up to a 1400 reduction in wages - equal to $\frac{1}{4}$ of a standard deviation decline. Appendix Table A7 has more detail on this result.

[27]The results reported are not marginal effects, so we cannot interpret the coefficient term as a marginal effect. Since this is a non-linear model, marginal effects are not possible to construct for interaction terms.

are historical homelands i.e. reservations created under Apartheid: Eastern Cape, KwaZulu-Natal, North West, Mpumalanga, and Limpopo. Job search costs here are very high and thus the chances of finding a job are low, due in large part to high transport costs. Therefore, I would not expect to find much change over the intensive margin. Regression discontinuity estimates for Bantustants are indeed statistically insignificant. I do not report them, but these results are available on request.

We can see a large, and statistically significant, increase in labor force participation over time for those living in Bantustans, indicating the end of the recession and the revival of the labor market. While the other outcomes do not suggest similar changes, they are all statistically insignificant. Taken together, these results imply labor demand picks up over time even for those in areas where employment prospects are poor.

Other Mechanisms The results so far suggest pensions affect labor supply through labor demand. Other studies have pointed out alternate ways in which pensions can affect labor supply. [Ardington et al., 2009] have argued the pension relaxes a child care constraint, allowing working age individuals - particularly mothers - to migrate for work. Further work [Ardington et al., 2016] argues 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 A8 but fail to find they are operational in explaining labor supply decisions. The effects are statistically insignificant when I consider mothers alone (Panel A of Appendix Table A8) or rural men alone (Panel B of Appendix Table A8).[28] While it would be preferable to further split the 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.

Related aspects 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 A9 has these results.

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. To see if a similar effect operates for labor supply, I restrict the sample such that there is no female pensioner in the 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 A10 has these results. Ambler [2016] argues such gender based differences come from a change in

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

relative bargaining powers within the household.

Abel [2019] notes that the effect of pensions is unclear for unemployed prime-aged adults. Pension resources can alleviate credit constraints thus increasing search but can also lead to an increase in reservation wages. I fail to reject a null effect of the pension on reservation wages (Appendix Table A11[29]). Under the mechanism I have described, pension resources affect labor supply only through a demand channel: if demand is weak, increasing search is less likely to yield any profitable result. Given the null effect on reservation wages, the mechanism I propose yields the following implication: the effect of the pension on unemployed prime-aged adults finding a job should become more positive over time. As the recession weakens, the benefits to increased search will rise, making employment more likely. I test for this indirect implication of the main hypothesis and find it holds: the probability of finding employment if previously unemployed is negative in the first two waves but becomes positive by the last wave. The effect of the pension is statistically insignificant in all these specifications; however, the confidence intervals do move rightwards (Appendix Table A12).

Blattman et al. [2014] suggests credit constraints matter, so when given cash transfers new businesses can start up. This implies labor supply should 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 or resident in house-

[29] 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?”

holds separate from the pensioners (Appendix Table A13).Therefore the low rate of enterpreneurship in South Africa cannot be due to credit constraints.

7 Conclusion

In this paper, I demonstrate pensions affect labor supply on the intensive margin through the demand side. Hours worked reduces by nearly 1.7 standard deviations in response to pensions - but only around the time when the national economy is in a recession, and only for workers in sectors affected by the recession. Pension arrival will increase the demand for leisure if it is a normal good as pension income raises household income. A recession lowers the opportunity cost of leisure, also increasing demand. Importantly, neither the pension nor the recession by themselves lower labor supply - it is the combination of the two events that appears to cause a reduction.

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, attention should be directly brought onto the labor market, if the concern is around how the pension program affects its unintended recipients. It isn't clear, for instance, why the rate of self-employment is so low in South Africa. Identifying the constraints to self employment can be useful as low self-employment could be drawing down labor mobility. High labor mobility can act as a buffer to episodes of tightness in labor markets.

It is unclear however whether the labor supply reduction - when it does take place - induced by the pension program is welfare reducing or welfare improving. Ultimately these 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. To the extent that labor supply is unaffected by the pension program, however, the deadweight loss would be minimal.

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.

Bertrand, M., S. Mullainathan, and D. Miller (2003). Public policy and extended families: Evidence from pensions in South Africa. *The World Bank Economic Review 17*(1), 27–50.

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.

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.

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.

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

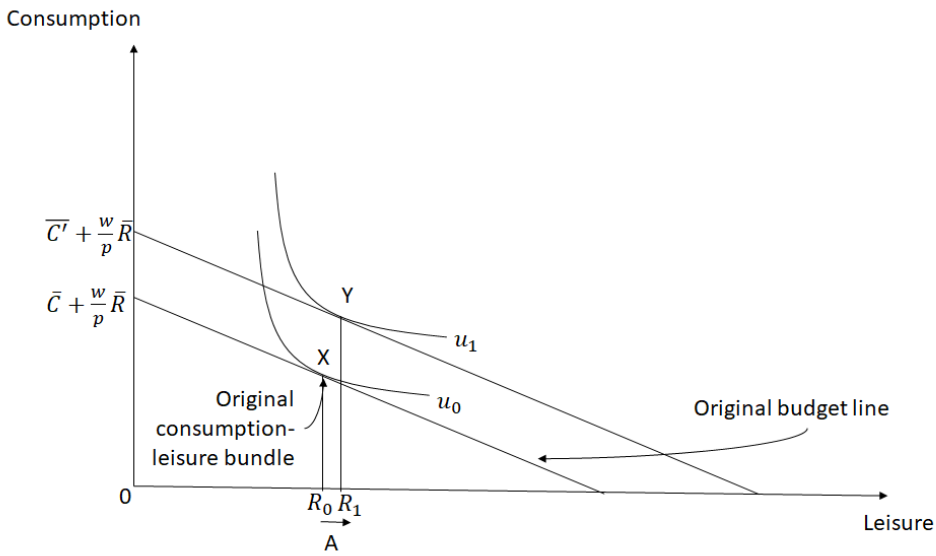

Figure 1: Pension Arrival Alone

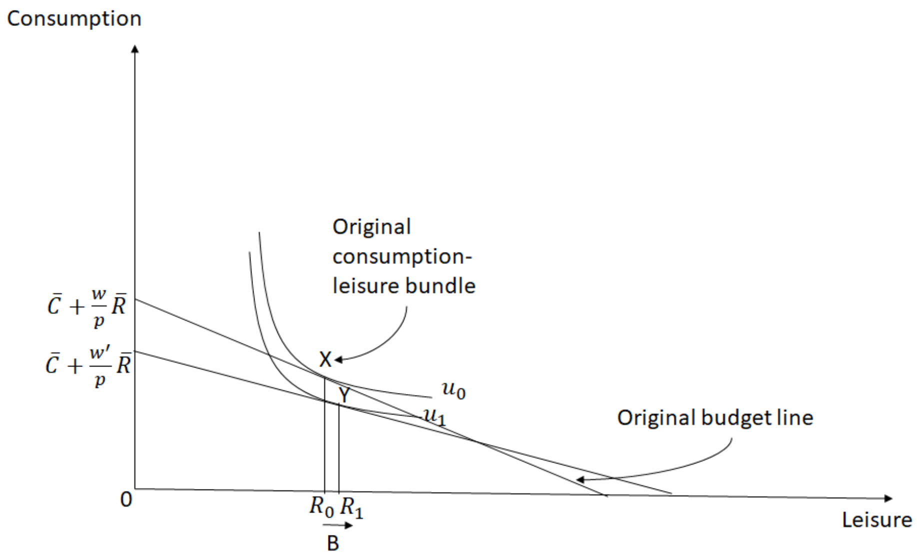

Figure 2: Recession Alone

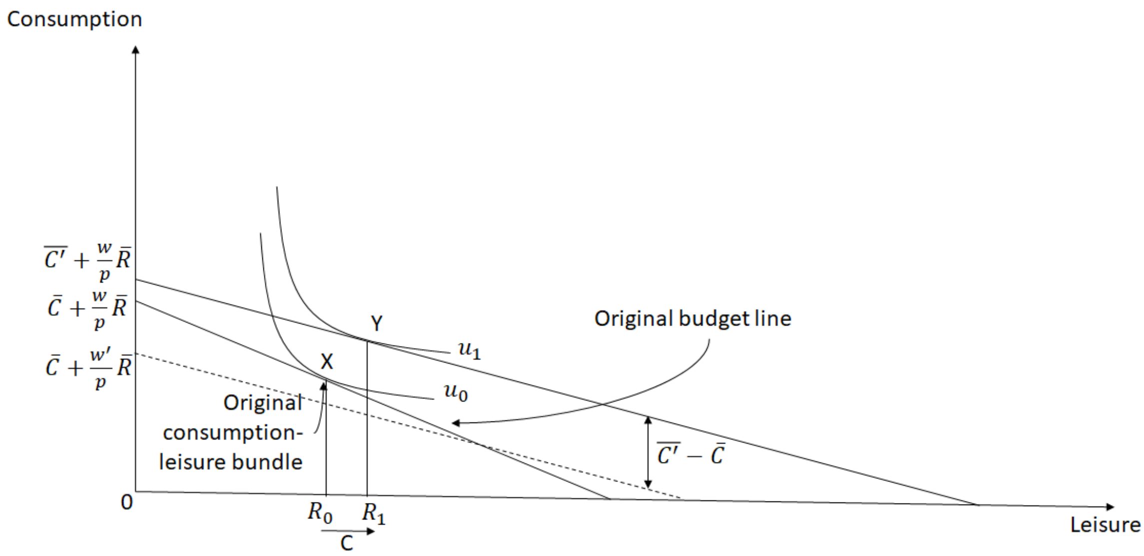

Figure 3: Pension Arrival × Recession

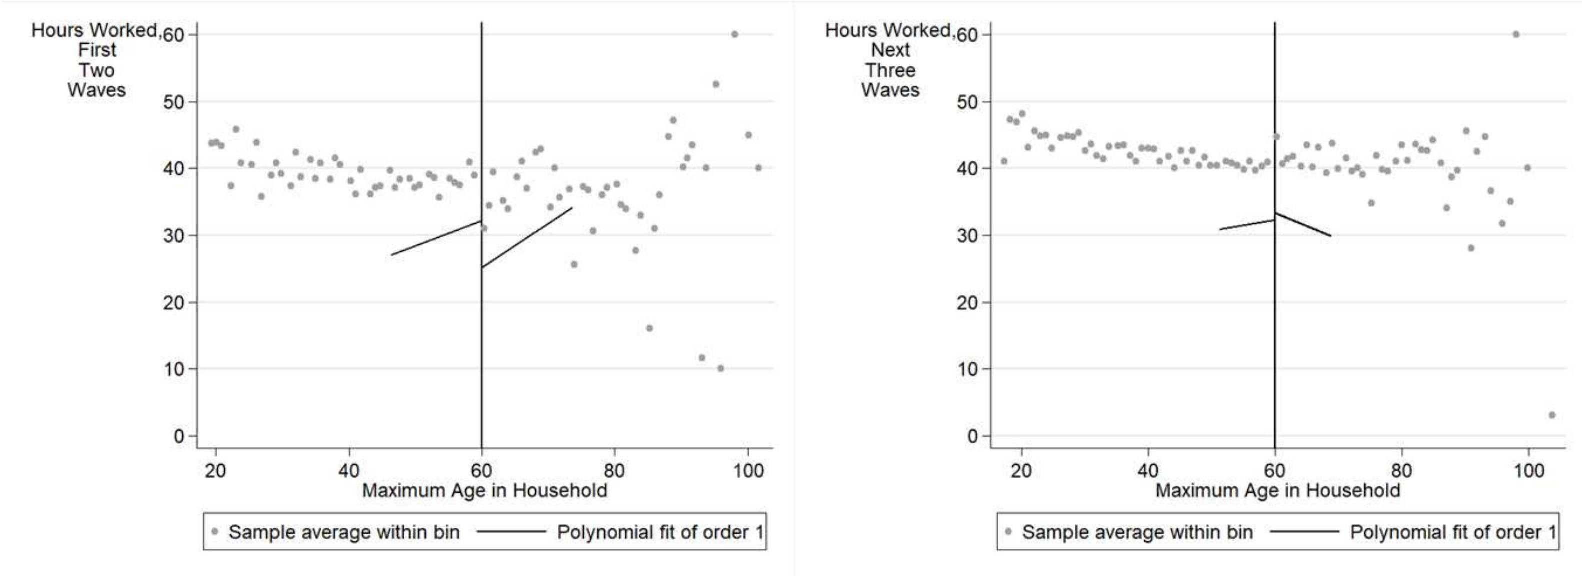

Figure 4: This figure shows the drop in hours worked - shown by the grey dots - once a household member reaches the pensionable age of 60, indicated by the vertical black line. The drop, however, occurs only in the first two waves. In the next three waves, no drop occurs; indeed, the hours worked nudge upward slightly. The figures also include a fitted regression line drawn across the bandwidth; the jump in these lines across the cut-off is used to derive the regression discontinuity estimate. Controls used to estimate the regression are gender, race, urbanization, household size, age and education. As indicated, a linear polynomial is used on both sides of the cut-off; and the bandwidth is chosen to minimize mean squared error. (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	
Waves	1 to 2	3 to 5	1 to 2	3 to 5
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	-	-	0.080	0.435
Eligible Males			(0.06)	(0.040)
Total Pension	-	-	0.658	0.615
Eligible Females			(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: Interpreting the Impact of the Pension Program on Labor Supply

Panel A: Household Composition Eliminated				
	(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 Estimate	-8.65	0.32	-25.67	-3.81
	(2.96)	(1.26)	(8.43)	(3.36)
Controls	Y	Y	Y	Y
Observations	4,502	9,868	904	2,430
Panel B: Sectors Hit by Recession				
	(1)	(2)	(3)	(4)
Sample	Recession Sector		Non-Recession Sector	
Period	2008-2011	2012-2017	2008-2011	2012-2017
Discontinuity Estimate	-20.80	-3.697	-39.17	8.388
	(9.76)	(5.79)	(24.68)	(7.28)
Controls	Y	Y	Y	Y
Observations	1,645	3,826	2,671	5,868
Panel C: Gender of Recipient				
	(1)	(2)	(3)	(4)
Sample	Male		Female	
Period	2008-2011	2012-2017	2008-2011	2012-2017
Discontinuity Estimate	-32.87	-1.449	-39.06	6.478
	(13.46)	(6.61)	(21.43)	(6.196)
Controls	Y	Y	Y	Y
Observations	2,188	4,774	2,314	5,094

Interpreting the Impact of the Pension Program on Labor Supply: Table 3 Continued

Panel C: Skill Levels

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

Panel D: Demand Side Estimates:

	(1)	(2)	(3)	(4)	(5)
Outcome	Upward Skill Mobility	Narrow Employed	Broad Employed	Discouraged	Participation
Wave 4 $\times$ Recession Sector	0.339				
	(0.169)				
Wave 4 $\times$ Bantustan		-0.094	-0.038	0.244	0.193
		(0.088)	(0.082)	(0.174)	(0.058)
Controls[b]	Yes	Yes	Yes	Yes	Yes
Observations	3,045	10,211	12,492	4,553	28,886
Individual FE	Yes	Yes	Yes	Yes	Yes
# Individuals	988	3152	3724	1399	7094

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 all discontinuity estimates is hours worked in the reported primary job, however for the results in columns (3) and (4) in Panel B the outcome is hours worked in all jobs. Self-employed individuals are included in the latter outcome.
Controls for Panels A through C are the same as in columns 1 and 2 in Table 2. *(b)*: Controls in Panel D are the same as in columns 3 and 4 in Table 2.

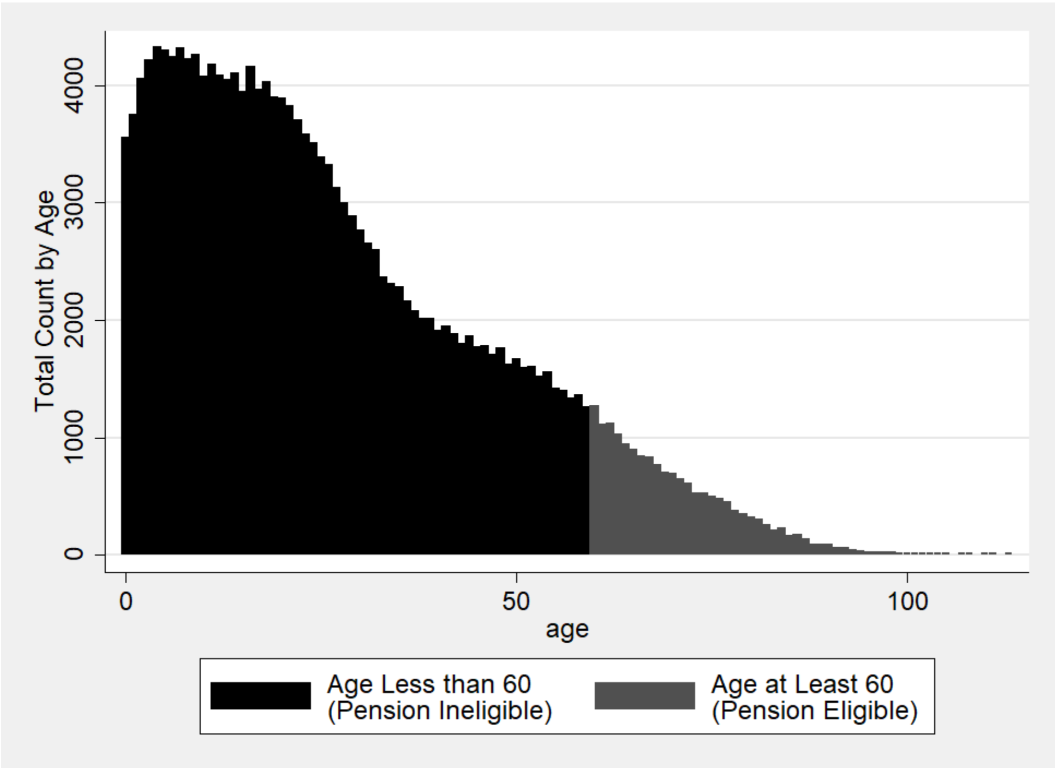

Figure A1: This figure shows the assignment to belonging to a pensioner household, as a function of the age of the individual. Each bar plots the total count of individuals of a particular age. The bars in dark grey show the counts for individuals below the (pension eligible) age of 60, and those in light grey show the counts for individuals above the age of 60. The lack of discontinuity at the pension eligible age of 60 indicates that the assignment to pension household is not being systematically manipulated. (Source: National Income Dynamics Survey Waves 1 to 5)

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)
Waves:	1 to 2	3 to 5	1 to 2	3 to 5	1 to 2	3 to 5	1 to 2	3 to 5
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:	1 to 5	1 to 2	3 to 5	1 to 5	1 to 2	3 to 5	1 to 5	1 to 2	3 to 5
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)
Waves:	1 to 5	1 to 2	3 to 5	1 to 5	1 to 2	3 to 5	1 to 5	1 to 2	3 to 5
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)					
Sample:	All Waves		Wave 1 and 2		Wave 3, 4 and 5	
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)					
Sample:	All Waves		Wave 1 and 2		Wave 3, 4 and 5	
First Stage Instruments						
Total Pension	0.350***	0.335***	0.269***	0.245***	0.366***	0.364***
Eligible Males	(0.015)	(0.016)	(0.033)	(0.034)	(0.023)	(0.023)
Total Pension	0.624***	0.572***	0.590***	0.535***	0.604***	0.571***
Eligible Females	(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: Transfers Between Households: Non-Resident Pensioners

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sample:	Only Non-Resident Pension Eligible				Only Non-Resident Pension Eligible Female			
Outcome:	Hours Worked Primary Job		Hours Worked All Jobs		Hours Worked Primary Job		Hours Worked All Jobs	
Waves:	1 to 2	3 to 5	1 to 2	3 to 5	1 to 2	3 to 5	1 to 2	3 to 5
Point Estimate†	-8.65	0.32	-4.60	0.32	-7.85	-2.07	-25.67	-3.81
Standard Error†	2.96	1.26	3.82	1.26	5.79	3.21	8.43	3.36
Lower 95% CI†	-14.46	-2.14	-12.08	-2.14	-19.20	-8.36	-42.20	-10.40
Upper 95% CI†	-2.85	2.78	2.89	2.78	3.50	4.23	-9.14	2.78
Outcome Mean	38.26	41.84	39.72	42.14	37.21	40.77	39.45	40.74
Outcome Standard Deviation	16.51	13.69	39.94	23.06	15.64	12.72	29.20	21.13
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Observations	4502	9868	6292	9868	680	1933	904	2430
Effective Observations								
Left of cut-off	833	2974	1221	2974	172	263	116	409
Right of cut-off	404	1477	598	1477	109	224	75	317

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 A7: The Wage-Hour Function

Hour	-508.6
	(168.1)
Hour#Wave = 2	-243.7
	(449.0)
Hour#Wave = 3	-679.2
	(243.1)
Hour#Wave = 4	-1,298
	(277.0)
Hour#Wave = 5	-1,436
	(314.5)
Controls †	Y
Observations	16,938
R-squared	0.168
Outcome Mean	3312
Outcome Standard Deviation	7330

The outcome variable is wages earned. 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.

†: Controls at the individual level include gender; race; 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. A constant is also included.

Table A8: 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		
Waves:	1 to 5	1 to 2	3 to 5	1 to 5	1 to 2	3 to 5
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 A8 Continued

Panel B: Labor Migration - Rural Men						
	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:	Hours Worked, Primary Job			Hours Worked, All Jobs		
Waves:	1 to 5	1 to 2	3 to 5	1 to 5	1 to 2	3 to 5
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 A9: Intergenerational Transfers Within Households

	(1)	(2)	(3)	(4)	(5)	(6)
Sample:	Overall		First Two Waves		Next Three Waves	
	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 A10: Gendered Impacts of the Pension on Labor Supply

	(1)	(2)	(3)	(4)
	At least One Female Pensioner		At least One Male Pensioner	
Waves:	1 to 2	3 to 5	1 to 2	3 to 5
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 A11: Individual Preferences: Reservation and Fair Wages

	(1)	(2)	(3)	(4)
Outcome	Reservation Wage		Fair Wage	
Sample	Wave 1 to 2	Wave 3 to 5	Wave 1 to 2	Wave 3 to 5
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 A12: Effect of the Pension on Finding a Job

	Outcome = 1 if newly employed, = 0 otherwise							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Wave $\leq$ 2		Wave > 2		Wave > 3		Wave > 4	
Household Pension		-0.022 (0.014)		-0.008 (0.012)		-0.006 (0.015)		0.024 (0.021)
First-Stage Estimates:								
Total Males $\geq$ 60 years	0.296 (0.023)		0.363 (0.019)		0.368 (0.021)		0.372 (0.029)	
Total Females $\geq$ 60 years	0.556 (0.020)		0.609 (0.014)		0.629 (0.017)		0.609 (0.021)	
Controls †	Y	Y	Y	Y	Y	Y	Y	Y
Observations	7,723	7,723	17,709	17,709	12,334	12,334	6,332	6,332
F statistic ‡	552.6	552.6	1583	1583	1320	1320	745.5	745.5
Hansen J Statistic	1.415	1.415	0.222	0.222	1.682	1.682	0.0950	0.0950
p-val	0.234	0.234	0.637	0.637	0.195	0.195	0.758	0.758
Outcome Mean	0.0365	0.0365	0.0862	0.0862	0.0967	0.0967	0.0855	0.0855
Outcome Standard Deviation	0.188	0.188	0.281	0.281	0.296	0.296	0.280	0.280

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.

†: 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. A constant term is also included.

‡: 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 A13: Effect of the Pension on the Self Employed

	(1)	(2)	(3)	(4)
	Co-Resident		Non-Resident	
	Wave ≤ 2	Wave > 2	Wave ≤ 2	Wave > 2
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.