Female Labor Supply and Rural Pension Eligibility in Brazil

Gaurav Khanna*
Margaret J. Lay†
Stephanie Lee‡
Benjamin Thompson§

July 11, 2024

Abstract

In 1991, Brazil expanded its rural old-age pension to cover millions of previously uncovered women, conditional on work requirements. We use a difference-in-differences approach to show that this expansion drastically increased women's employment by nine percentage points, or 26 percent. This increase in labor force participation occurred among women who were immediately age-eligible, and among younger cohorts that would be eligible in the future. These results illuminate the capacity of workers to respond to financial incentives for labor participation in old age, and the extent to which younger workers might be forward-looking as they respond to retirement incentives.

JEL: H55, J26, O54

Keywords: Public pensions, retirement, female labor supply

*University of California – San Diego
†Mount Holyoke College
‡University of Washington
§Independent Researcher

We thank Breno Braga, Jim Hartley, Shanthi Ramnath, and Emily Conover for insightful discussions, and seminar audiences at various conferences for helpful suggestions.

1 Introduction

Labor force participation among women in rural Brazil increased by a surprising 26 percent between 1990 and 1992. This dramatic growth in female labor supply, one of the world's largest seen in a span of a few years, occurred on the heels of a generous expansion of the rural pension regime. Pension expansions are usually found to depress labor supply through a negative wealth effect (Huang & Zhang, 2021; Bando et al., 2016, 2020, 2022; Kaushal, 2014; de Carvalho Filho, 2008). This expansion, however, created a competing eligibility incentive that, we show, was particularly effective in increasing labor-force participation for women likely because they were initially less attached to the labor force. In this paper, we model the various labor-supply incentives created by this 1991 rural pension expansion, and use annual large-scale household data and a difference-in-differences specification to show that the pension led to a dramatic increase in women's labor supply on the extensive margin. The design of transfer systems can influence women's labor-force participation more broadly (Kleven, 2019; Bastian, 2020). Understanding how transfer design facilitates female labor supply has particular implications for economic development, as increases in female labor-force participation are associated with prosperity and broader improvements in gender relations (Goldin, 1995; Dinkelman & Ngai, 2021; Blau & Kahn, 2013; Anderson & Eswaran, 2008).

The rural pension expansion, passed and implemented in 1991, newly provided a non-contributory pension to women aged 55 and older who were not household heads, provided that they could produce evidence of previous rural work. Our extended difference-in-differences specification compares the pension receipt and labor supply of married rural women to that of married urban women (first difference), before and after the reform (second difference), to find a sustained increase in labor supply among married rural women of nine percentage points, or approximately 26 percent. This increase was immediate among all cohorts, but larger and short-lived among women who were near retirement age when the expansion took effect, and smaller but sustained among women in younger cohorts. These findings suggest that women who might not otherwise enter the labor force adjust their labor supply when the pension incentive is strong enough. Older women will work to gain eligibility, and younger women will increase their labor supply in anticipation.

Our finding that an increase in pension generosity is associated with an increase in labor supply is uncommon among the literature exploring old-age pensions and labor supply. Much of the existing literature documents a wealth effect: more pension generosity decreases labor supply (Huang & Zhang, 2021; Bando et al., 2016, 2020, 2022; Kaushal, 2014; de Carvalho Filho, 2008), while less pension generosity increases labor supply (Staubli & Zweimüller, 2013; Neumark & Song, 2013; Brown, 2013; Mastrobuoni, 2009; Duque, 2021). Some papers identify this wealth effect as the difference in labor supply between age-eligible individuals and non age-eligible individuals (Bando et al.,

2016; Duque, 2021; Shu, 2018; de Carvalho Filho, 2008). However, that difference is confounded by labor supply responses among the non-age eligible as the non-age eligible may also change their lifetime labor supply. Our empirical specifications avoid the need to use slightly younger cohorts to control for time-specific effects by comparing the impact of the pension expansion on female labor supply to its impact on various other similarly-aged control groups. Further, and perhaps more importantly for explaining why our results suggest an opposite sign from previous literature, the reform studied here creates an incentive to increase labor supply to achieve pension eligibility, in addition to the traditional wealth effect that discourages labor supply. Our results suggest that this eligibility effect is particularly powerful in bringing married women into the workforce.

A second robust finding in the literature is that people tend to retire at discontinuously higher rates when they reach the age of pension eligibility (Neumark & Song, 2013; Behaghel & Blau, 2012; Shu, 2018). Previous work has emphasized credit constraints, discontinuous marginal labor supply incentives, and reference-dependent preferences at pension eligibility ages to explain this behavior (Seibold, 2021; Lalive et al., 2023; Gruber et al., 2022). We next use a difference-in-discontinuities specification to explore whether women living in rural areas in Brazil make labor supply choices that replicate this pattern. We find little evidence of a discontinuous decrease in labor supply at 55, the age at which women who work in agriculture become eligible for the rural pension, immediately after the reform. However, this labor supply response develops among younger cohorts of women: women who turn 55 in 2006, for example, are five percentage points less likely to work than they are at the marginally younger 54. These findings suggest that women who are aware of the pension eligibility age and work requirements at the beginning of their working life exhibit a discontinuous decrease in labor supply at the age of eligibility.

The implications of our results are nuanced in a context like Brazil, where only ten percent of rural households had a woman as their household head prior to the reform. An earlier pension scheme was targeted at heads, implying that mostly men had access to public pensions. As a result, women's total compensation for the same workload was lower than men's, discouraging female labor force participation. If, additionally, the prevailing market wage adjusts downward to account for this publicly-provided compensation to (mostly male) workers, highly-productive women may stay out of the market; even as they would have participated in the absence of these cultural and economic frictions. The de-facto targeting of the 1991 pension expansion to women may have helped alleviate such frictions and driven our documented increase in female labor force participation.

We explore these various competing effects of the reform on female labor supply by building a model of labor supply decisions over the life cycle, for people with heterogeneous utility of home- relative to market-produced goods. This model suggests that there are four groups of people driven by countervailing effects – a characterization of labor supply behavior that echos the bunching behavior documented in other contexts (Brown,

2013; Manoli & Weber, 2016; Seibold, 2021). First, those who receive high utility from home relative to market production are unlikely to increase labor-force participation and, instead, forgo accessing the pension. Second, a group of marginal workers will now work more years so as to meet the work eligibility requirement and obtain the pension. Third, a group of marginal workers who would have worked more than the work requirement without the expansion, reduce their labor supply to the required amount. Finally, a group of workers that would have worked many more years than required by the expansion instead reduce their labor supply somewhat (and still work more than required) given their extra expected pension wealth. We use this model to compare lifetime labor supply decisions with and without the pension, and explore the transition response to the pension expansion among Brazilian women with a cohort-level smoothing assumption.

The model allows us to document competing effects in labor supply decisions, and quantify the aggregate economy-wide changes to cohort-specific work participation. Together with our empirical patterns, the model sheds light on aggregate lifetime labor supply, differential responses to the pension expansion by cohort and gender, and adjustments in part-time work. The model guides us in understanding how wealth, eligibility, and retirement-timing effects of the expansion influence labor supply over the lifetime, given the treatment effects we estimate. The standard wealth effect of expanded pension benefits encourages people to work less over their lifetimes, while the work requirement creates a separate eligibility effect encouraging some to work more over their lifetime. The retirement-timing effect may be positive or negative in any given year, but sums to zero over the lifetime of a cohort, as the cohort's target retirement age adjusts to new pension incentives.

We find that, as a result of the pension, lifetime labor supply increased by between 6.5 percent (for younger women) and 10.5 percent (for older women), indicating that the eligibility effect was particularly strong for older women. Part-time work among rural married women increased by 20 percentage points, with much of that increase due to new labor market entrants. The female labor-force participation rate increased from 37 percent to 57 percent of the male labor-force participation rate.

Our theoretical approach to interpreting labor supply responses to pension incentives differs from the option-value models standard in the retirement literature (Stock & Wise, 1990; Samwick, 1998; Coile & Gruber, 2007). Those models focus on discontinuous returns to working an additional year when an individual is near a retirement eligibility cutoff, and are critical in understanding marginal labor supply decisions near retirement age of groups that are strongly attached to the labor force. Our approach allows us to focus on the impact of the pension expansion, throughout the working life, on groups that are unlikely to work without the expansion. In this way, it is reminiscent of the literature debating the impact of the US's Earned Income Tax Credit and other conditional transfer programs on women's extensive- and intensive-margin labor supply decisions (Kleven,

2019; Whitmore Schanzenbach & Strain, 2021; Verho et al., 2022). This theoretical focus motivates an empirical approach that differs from those often used to explore the labor supply impact of old-age pension programs in developing countries (Bando et al., 2016, 2020, 2022; Kaushal, 2014). Rather than comparing the age-eligible to non-age eligible – a comparison that necessarily highlights negative wealth effects or reference-dependent retirements – our focus on lifetime labor supply reveals nuanced behavioral responses in which expansions in generosity may increase labor market participation.

2 Background

The Brazilian Constitution of 1988 initiated a dramatic expansion of old-age pensions in rural areas that had a particularly large impact on women. The rural pension system in place prior to this reform, established in 1971 and referred to as PRORURAL, granted an old-age pension equal to 50% of the minimum wage to the head of all rural households upon turning 65, provided that the household head produced evidence of working in the rural sector in one of the previous three years. As heads of households were primarily men, most married women were not eligible. Receipt of the rural pension was not means- or retirement-tested. A separate social security system covered Brazilians living in urban areas, in which both men and women, regardless of whether they headed their household, were eligible to receive a pension at age 70, or after 30 years of work. The urban pension amount depended on a recipient's past years of work and recent labor earnings, but was bounded below by 90% of the minimum wage. Receipt of the urban pension required recipients to quit their current job, though they could continue working elsewhere.

The 1988 Brazilian Constitution committed to equalizing this discrepancy in rural and urban pensions. Law (*Lei*) #8212/8213, passed in 1991, stipulated the details by which that equality would be achieved (see Table A.1). This law made minor changes to the urban pension scheme and substantial changes to the rural pension scheme. The law adjusted the urban pension scheme by increasing the minimum benefit amount to 100% of the minimum wage, initiating a tax on covered wages, removing the requirement that recipients quit their current job, and decreasing the work requirement to 25 years for women. In the rural pension scheme, *Lei* #8212/8213 expanded access to old-age pensions to household members other than the household head; increased the benefit amount to 100% of the minimum wage; and reduced the eligibility age from 65 for all recipients, to 60 for men and 55 for women. Further, the law increased the number of years of work required for pension eligibility: to receive the pension in 1991, rural individuals were required to produce evidence that they worked in a rural occupation for at least 5 years, though those years of work could be discontinuous and anecdotal evidence from local pension administrators suggest that there was a low bar for what constituted evidence of rural work. This work requirement increased gradually for subsequent cohorts, so that people who attained the age of pension eligibility in 2011 or later were required to have

worked for at least 15 years in a rural occupation. As a result, millions of married women were newly eligible to receive an old-age pension at age 55, provided that they could produce evidence of rural work history.

This reform had different impacts on work incentives for rural men and women. Most rural male pension recipients had worked more than 15 years prior to the expansion, so the newly expanded work requirement was not binding for most men. As a result, the main impact of the pension expansion on rural men was to double the size of the pension received and decrease the eligibility age from 65 to 60. This particularly increased the value of future pension benefits for any man below age 65, and the value of current pension benefits for newly eligible men between ages 60-65. de Carvalho Filho (2008) uses a triple-difference approach to assess the impact of this positive income effect on the labor supply of newly eligible men. By comparing rural to urban men (first difference) and men aged 60-64 to those aged 55-59 or 65-69 (second difference), before and after the reform (third difference), they find that newly age-eligible men decreased employment by 38 percentage points.

For married women in rural areas, on the other hand, the minimum rural work requirement established with *Lei* #8212/8213 was more likely to be binding. Prior to the reform in 1987, only 36 percent of rural women aged 25-75 worked, while 91 percent of rural men did so. Unlike men, for whom the 1991 pension expansion primarily increased lifetime wealth and thus exerted negative pressure on labor supply, newly eligible women who were not considered household heads faced an incentive both to decrease labor supply due to an increase in lifetime wealth, and to increase labor supply to attain pension eligibility. A cursory review of employment patterns among rural women shows that both the age-eligible and age-ineligible women increased their labor supply in 1991, suggesting that the latter eligibility incentive dominated the labor supply decisions for many women. The following sections explore this empirical result in detail.

3 Data and Descriptive Statistics

The *Pesquisa Nacional por Amostra de Domicilios*, or PNAD, is an annual, cross-sectional survey of approximately 100,000 households that began in 1967. The survey emphasizes labor-market activity and is representative of the Brazilian population, aged 14 and above. It asks detailed questions about work and demographic aspects of household members. Importantly for this study, it contains information on household members' pension receipt and work status, including work in informal employment. Until 2004, the survey included urban areas in all states of Brazil and rural areas in all states except those in the northern region. Rural areas in the northern region were included beginning in 2004. This paper considers adults between ages 25 and 74, and uses data compiled between 1981-2013, omitting years in which the PNAD was not conducted. Our main analysis excludes the northern region, but our results are robust to including that region.

Table 1: Characteristics of Rural and Urban Women and Men, before and after the reform

	Women				Men			
	Rural		Urban		Rural		Urban	
	Before	After	Before	After	Before	After	Before	After
A. Pension and Labor Force Status								
% identifying as HH head	12	17	20	32	89	83	87	76
% of population receiving pension	7.2	18	7.8	11	11	15	14	14
% of population 55+ receiving pension	28	66	30	40	40	55	56	57
% of population 65+ receiving pension	52	79	44	55	78	88	83	85
% living in HH receiving ≥ 1 pension	3.2	13	4.0	6.0	2.5	12	3.4	5.8
% worked in reference week	34	50	40	51	91	88	82	79
Average hours worked per week	13	15	16	19	45	39	39	36
B. Individual and Household Characteristics								
Average age	43	44	42	44	43	44	42	43
% married	79	77	66	62	82	76	81	74
Avg. number of children in HH	2.5	1.5	1.8	1.1	2.4	1.5	1.8	1.1
Avg. number of adults in HH	2.8	2.7	2.9	2.7	2.8	2.7	2.9	2.8
% living in multigenerational HH	24	25	29	28	20	21	24	24
N Observations	146,205	224,723	563,653	1,412,368	155,452	242,644	499,368	1,225,164

Notes. The sample contains all PNAD respondents between ages 25 and 75, from the years 1981 through 2013, excluding 1983, 1991, 1994, 1996, 2000, and 2010 as the PNAD was not conducted in those years. A respondent is classified as rural if they live in a rural village, and urban otherwise. Columns labeled "Before" include years between 1981 and 1990; columns labeled "After" include years between 1992 and 2013. "HH" refers to the household, or family unit in the PNAD. The average hours worked per week refers to the average working hours among the joint working and non-working populations.

Table 1, Panel A describes the pension and labor force status of rural and urban residents in the PNAD, aged 25 to 74, between 1981 and 1991 (before the reform) and between 1993 and 2013 (after the reform). Women are substantially less likely than men to identify as the household head in both rural and urban areas. The pension reform that expanded eligibility to non-household heads was associated with an increase in pension receipt among rural residents: from 28 to 66 percent among women aged 55 and older, and from 40 to 55 percent among men aged 55 and older. Following the reform, 13 percent of rural women and 12 percent of rural men live in households receiving multiple pensions; an increase from before the reform. Table 1, Panel B describes individual and household characteristics. The number of adults living in a household is comparable among rural and urban women, but rural women are more likely to be married and less likely to live in a multi-generational household than their urban counterparts. Consistent with Danzer & Zyska (2022), the number of children living in the household decreases slightly more in rural areas than in urban areas following the reform.

4 Empirical Results

These descriptive statistics show a dramatic expansion of the rural female labor force in Brazil, following the pension expansion in 1991. This section uses descriptive trends, a differences-in-differences specification, and a differences-in-discontinuity specification to examine this increase. The patterns indicate that the rural pension expansion led to a sizeable increase in female employment and brought married women into the workforce.

4.1 Trends in Pension Receipt and Labor Supply

We begin with an analysis of aggregate trends in married women's labor force participation by rural versus urban status, before going into a difference-in-differences analysis. As difference-in-differences estimates may hide which group of individuals are driving the changes, these simple descriptive trends transparently show that pension eligibility increases labor force participation among married women in rural areas.

Figure 1 describes patterns of pension receipt and labor supply among married women living in rural and urban areas in Brazil from 1981 through 2010. The left-most figure in Panel A shows that, while there was an expansion in pension receipt among married women living in both rural (dark) and urban (light) areas throughout this period, the expansion among married women in rural areas was particularly pronounced between 1990 and 1993 (Figure 1, Panel A).[1] The remaining figures in Panel A show that the rural workforce increased dramatically on both the extensive margin and in aggregate over that two-year period: the fraction of married women in rural areas that worked increased by 13 percentage points (37 percent) and the average length of the workweek among all

[1]This is consistent with de Carvalho Filho (2008), who shows that the Brazilian government took roughly two years to expand the rural pension system to reach the newly eligible.

Figure 1: Women's Pension and Work Status in Rural and Urban Areas by Marital Status

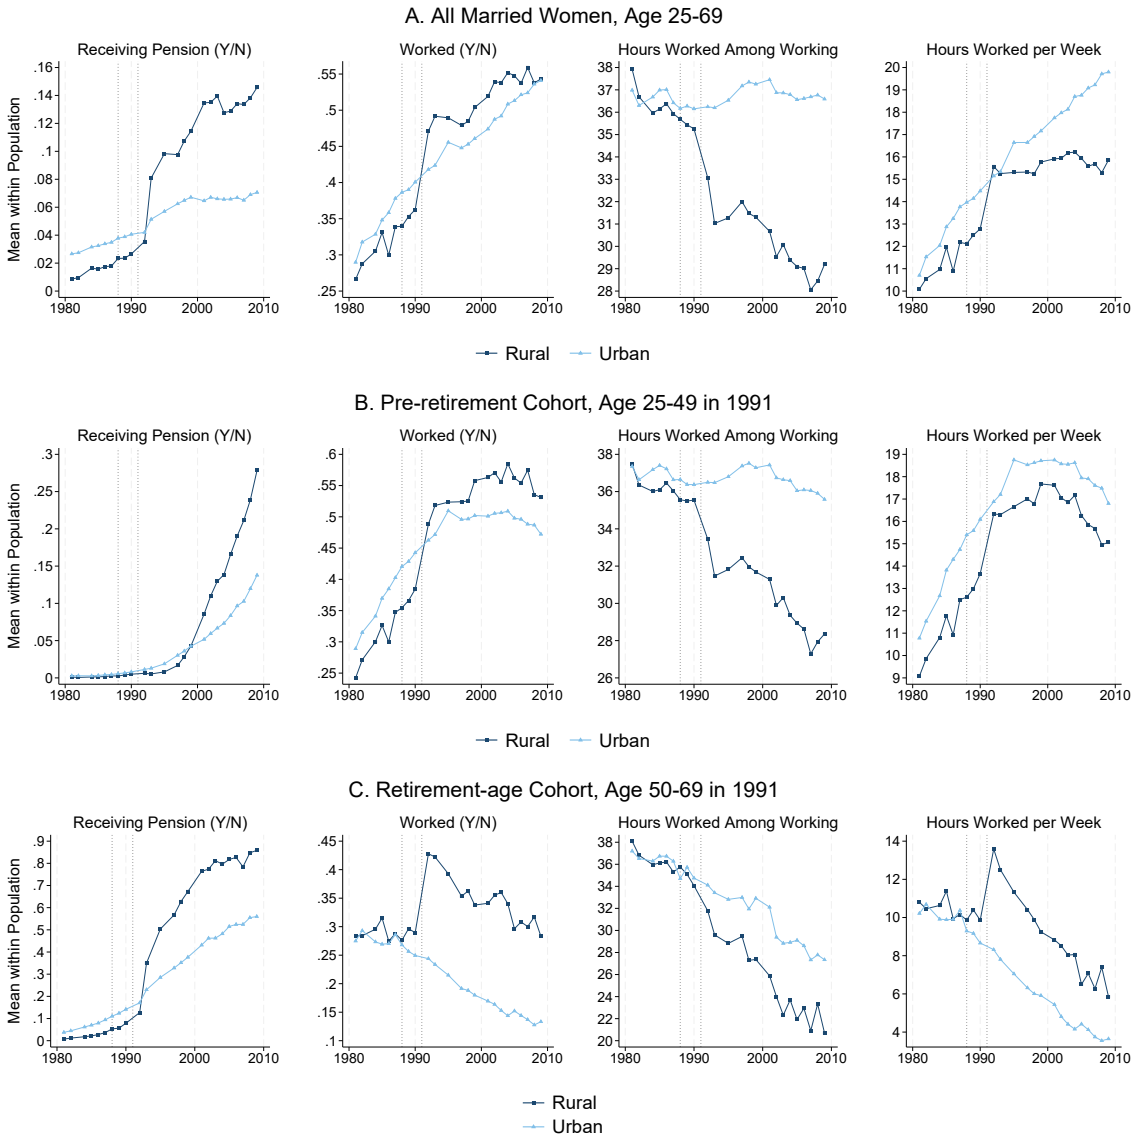

Notes. Panel A shows pension and labor force status among married women, ages 25-74, in rural areas (dark blue lines) and urban areas (light blue lines) from 1981 through 2013. Panel B shows pension and labor force status among married women who were between ages 25 and 49 when the law was passed in 1991, in rural areas (dark blue lines) and urban areas (light blue lines). Panel C shows pension and labor force status among married women who were between ages 50 and 69 when the law passed in 1991, in rural areas (dark blue lines) and urban areas (light blue lines). "Pension (Y/N)" refers to the fraction of the population that received a pension in each year. "Worked (Y/N)" refers to the fraction of the population that worked in the reference week in each year. "Hours worked among working" refers to the average number of hours worked per week among people who worked in the reference week in each year. "Hours worked per week" refers to the average hours worked per person among the full population, working and not working, in each year.

married women in rural areas increased by 2.5 hours (20 percent). The average length of the workweek among married women who worked, however, declined by 4.2 hours (12 percent) between 1991 and 1993. Labor supply remained elevated on the extensive margin

among married women in rural areas in the decades following the pension expansion, but aggregate hours worked flattened out shortly after 1993. This aggregate trend was influenced by a steady decrease along the intensive margin, measured as hours worked among the working, throughout the 2000s. These trends are particularly pronounced in comparison to labor supply trends among married women in urban areas, who experienced steady aggregate growth in labor supply, with steady increases along the extensive margin and very little change on the intensive margin.

Figure 1, Panels B and C, describe how different cohorts of married women adjusted their labor supply following the rural pension expansion. Panel B includes only married women who were between 25 and 49 in 1991 (pre-retirement age) and Panel C includes only women who were between 50 and 69 (near-retirement age) in 1991. Rural married women who were younger than retirement age increased employment on the extensive margin by 13.5 percentage points (35 percent) between 1990 and 1993, and sustained this increase throughout most of their careers. Rural married women close to or in retirement, between ages 50 and 69 in 1991, similarly increased labor supply by around 13 percentage points; this large peak was short-lived, but married retirement-age women in rural areas continued to work more than their counterparts in urban areas for many years following the reform. Among both cohorts, the growth on the extensive margin was short-lived, but the decline on the intensive margin continued throughout the 2000s. These patterns are again particularly pronounced in comparison to labor supply trends among married women in urban areas.

4.2 Difference-in-differences Specifications

We estimate the average annual treatment effect of the pension expansion with a difference-in-differences specification. We compare labor supply outcomes of people living in rural areas to those of people living in urban areas, over time. The following equation describes the extended difference-in-differences specification:

$$y_{irst} = \alpha \text{Rural}_r + \sum_{j=1981,\neq 1987}^{2013} \beta_j \text{Rural}_r + \delta_t + \mu_s + \Gamma'_{ist} X_{irst} \tag{1}$$

The outcome variable of interest, y_{irst}, is measured for individual i living in geographical area r of state s in year t. The treatment variable, Rural_r, is equal to one for individuals living in rural areas who could benefit from the newly expanded rural pension system, and zero for individuals living in urban areas. The coefficients δ_t and μ_s represent year and state fixed-effects. The vector of controls, X_{ist}, is included in robustness checks described in Section 4.2.2. These include controls for state-by-year fixed effects, state-by-rural fixed effects, age, and household demographics. We run this specification on various samples of people between the ages of 25 and 69, in the years 1981 to 2013. The omitted year is the year 1987, immediately before the constitutional reform.

4.2.1 Difference-in-differences: Baseline Results

Urban vs Rural Married Women. Our main estimates focus on married women living in rural areas. Figure 2, Panel A presents the coefficient estimates and 95% confidence intervals found by running equation 1 on the sample of married women in rural and urban areas, using pension receipt and three measures of labor supply as the outcome variables. The left-most figure of Panel A shows that pension receipt among married women in rural areas increased by four percentage points relative to their urban counterparts within two years of the pension expansion, and by nine percentage points relative to their urban counterparts by 2010. The fraction of married women in rural areas who worked, shown in the next figure of Panel A, similarly increased relative to its urban counterpart immediately after the reform, by nine percentage points (26 percent) between 1991 through 1993, and remained high until 2009. The average length of the workweek among working married women in rural areas, however, declined by two hours in 1992, by six hours in 2009, and remained at this low level through 2010. The final figure in Panel A shows that the increase in labor supply along the extensive margin dominated the decrease along the intensive margin in the early years: the average length of the workweek among all women increased by approximately two hours (12 percent) from 1991 through 1993. However, the overall treatment effect fell to zero within five years, and below zero by 2007.

Figure 2, Panels B and C present the primary estimates from specification 1 separately for women in pre-retirement and retirement-age cohorts. As expected, pension receipt increases quickly among age-eligible married rural women. The largest increases in labor supply along the extensive margin are among married women who are near age-eligible for the pension. Married women in rural areas who were between 55 and 69 were 20 percentage points more likely to work in the year following the enactment of Law #8212/8213. The increase in labor supply for younger cohorts, on the order of five percentage points for those who were between 25 and 44, and 15 percentage points for those between 45 and 54, in 1991 is smaller initially but more persistent than that for older cohorts. Despite the decline in the average length of the workday for most cohorts, overall labor supply increased for middle-age cohorts, between 45 and 64 in 1991, throughout the period considered.

Comparing Men and Women. While our focus is on married women's labor force participation, we may also expect married men to respond to aspects of the rural pension expansion. Men who were expecting to draw on the rural pension prior to the reform saw their eligibility age decline by five years and their benefits increase from 50 percent to 100 percent of the minimum wage. Figure 3, Panel A, compares descriptive trends in pension receipt, work status, and hours worked for married men to those of married women, in urban and rural areas. These figures confirm that the starkest increase in pension receipt

Figure 2: Pension and Work Status in Rural versus Urban Areas among Married Women, Difference-in-Difference Estimates

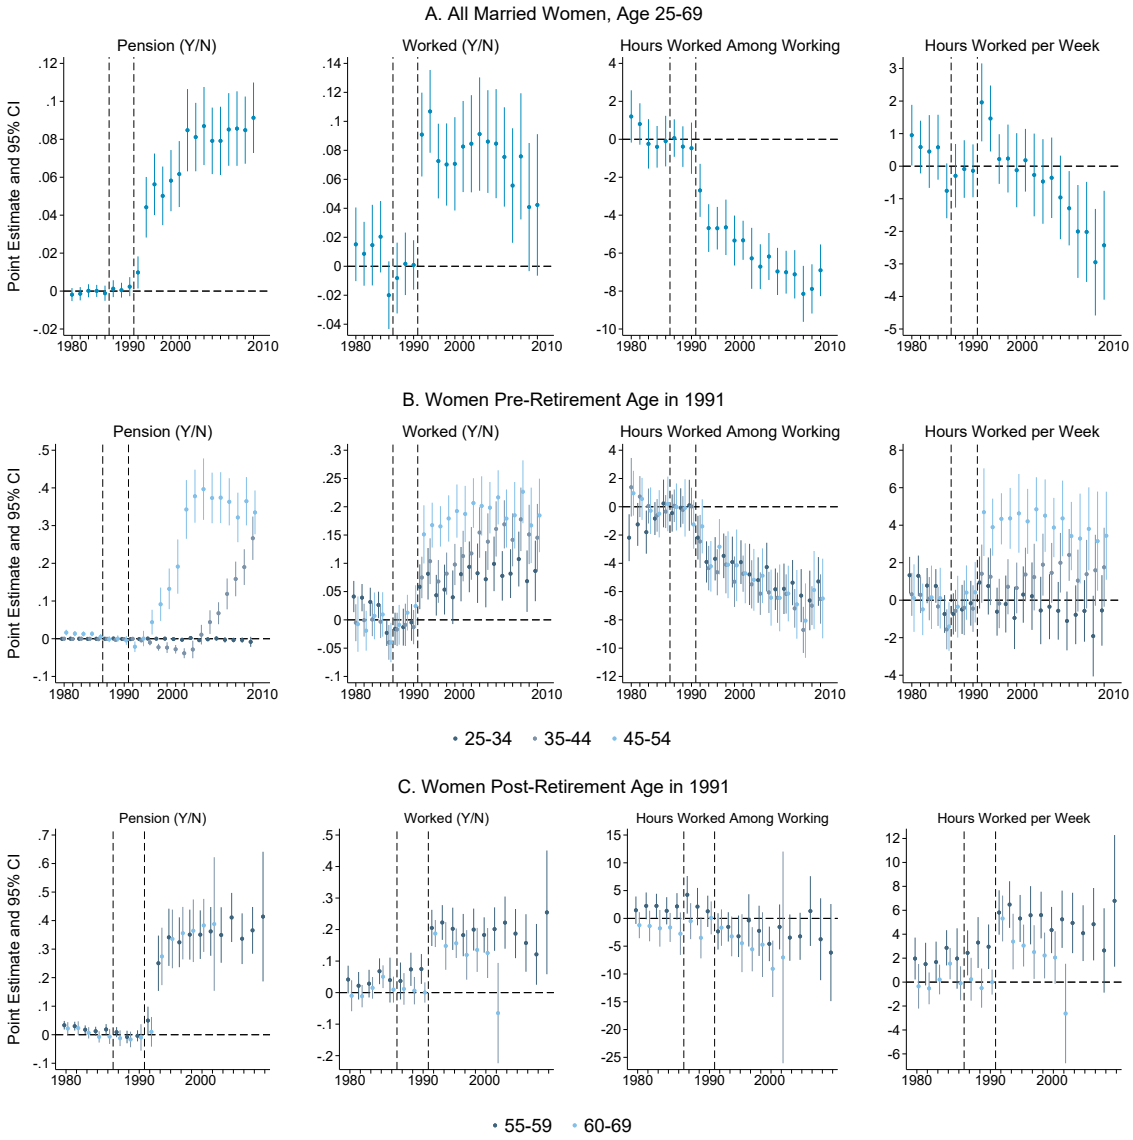

Notes. Each panel shows the β coefficient estimates and 95% confidence intervals on each year from an extended difference-in-difference regression of the form $y_{ist} = \alpha \times \text{Rural}_{ist} + \sum_{j=1981}^{1988} \beta_j^{pre} \times \text{Rural}_{isj} + \sum_{j=1989}^{2013} \beta_j^{post} \times \text{Rural}_{isj} + \delta_t + \mu_s$, where y_{it} is the outcome variable of interest and Rural = 1 if the individual lives in a rural area. Panel A includes all married women aged 25-69 within the year plotted. Panel B includes three different cohorts of married women who were younger than the pension eligibility age of 55 when the law was passed in 1991. Panel C includes two different cohorts of married women who were older than the pension eligibility age of 55 when the law passed in 1991. All panels are limited to women aged 25-69 within the year plotted. Coefficients are estimated relative to 1987, the year before the constitutional amendment announcing expansion of the rural pension scheme; the dashed vertical lines represent 1987 and 1991, the year in which the implementation of the expansion was announced. The title of each graph refers to the outcome variable used to generate the difference-in-differences estimates. "Receiving Pension (Y/N)" is an indicator variable equal to one if the individual is receiving pension payments. "Worked (Y/N)" is an indicator variable equal to one if the individual reports working in the reference week. "Hours Worked among Working" is a continuous variable equal to hours worked per week; this sample is limited to individuals who worked in the reference week. "Hours worked per week" is a continuous variable equal to hours worked per week; the sample includes all working and non-working individuals.

was among married women in rural areas immediately following the reform. However, pension receipt also increased discontinuously for married men in rural areas in that time period. The descriptive trends in labor supply measures, however, suggest that the largest adjustments in labor supply immediately following the reform were among married women in rural areas while longer-term declines in labor supply also occurred among married men in rural areas.

Figure 3, Panel B, compares these patterns more formally by presenting the difference-in-difference estimates on married women between the ages of 25 and 69 living in rural versus urban areas, first shown in Panel A of Figure 2, beside those found by running that same specification on married men of the same age in rural versus urban areas. The first graph in Panel B shows that pension receipt did, indeed, increase among married men in rural relative to urban areas. This increase is about half as large as the analogous increase for married women. There was also a small increase in labor market participation among rural married men, on the order of two percentage points, and a decline in average hours worked among working married men following the implementation of the reform. The long-term decline in hours worked per week was similar among married women and men.

We next explore this behavioral difference between men and women further by expanding specification 1 into a triple-difference that compares married women to married men, in rural versus urban areas, before and after the reform. Figure 3, Panel C, presents the triple-difference coefficients from this estimation. The first graph shows that rural married women experienced a two to four percentage point larger increase in pension receipt than rural married men, following the reform. Rural married women's subsequent increase in labor market participation (second graph) was four to six percentage points larger in magnitude, and their subsequent decrease in the average length of the workday (third graph) was two to five hours larger in magnitude, than rural married men. The flattening out of the triple-difference estimates between these two groups in later years suggests that the downward turn in labor supply that began in rural areas around 2007 was the result of some rural versus urban trend rather than the pension expansion itself.

4.2.2 Difference-in-differences: Identification and Robustness

Two identifying assumptions underlie this specification. First, the parallel trends assumption requires that rural and urban labor supply would have similar trends in the absence of the expansion in the rural pension system. The point estimates of the pre-trend coefficients, $\beta_j \ \forall j \in [1981, 1987]$, are not far from nor statistically different from zero, which somewhat alleviates the concern that this assumption is violated. Second, the exogeneity assumption requires that no other changes occurred simultaneously with the pension reform in 1991, besides the policy change of interest, that influence rural and urban labor supply choices in different ways. Under these identifying assumptions,

Figure 3: Comparison of treatment effects: Married women and Married Men

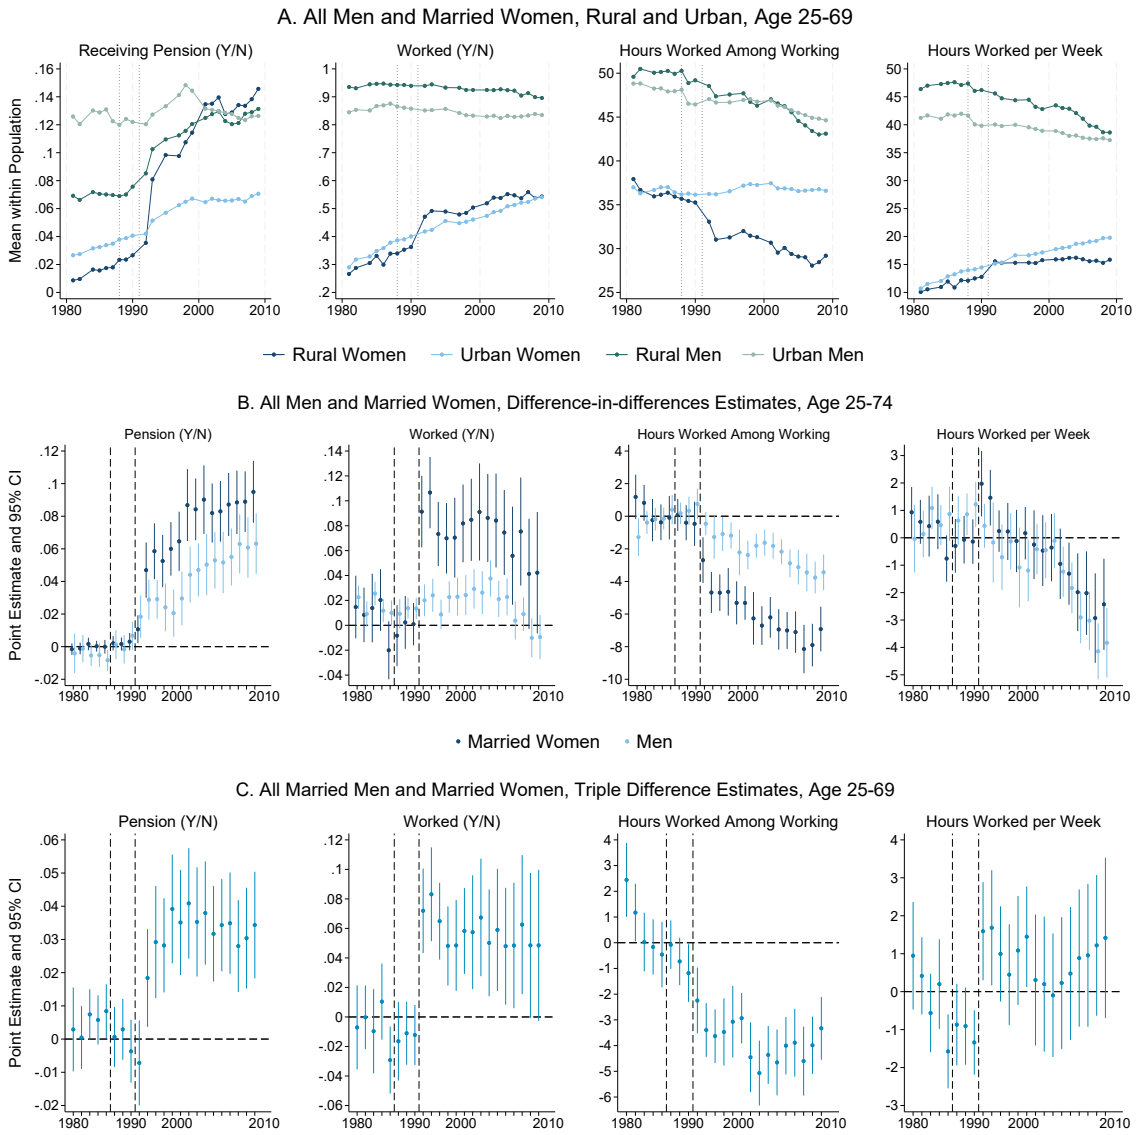

Notes. This figure compares the impact of the 1991 expansion on married women to that on two other groups of workers that had access to the rural pension prior to the reform – men and single women. Panel A compares the descriptive patterns in pension receipt and labor supply for men to those for married women (first shown in Figure 1.A.). Panel B compares the difference-in-differences estimates derived by limiting the regression sample to Men, to the main difference-in-differences estimates derived from limiting the regression sample to married women (first shown in Figure 2.A.). Panel C presents the triple-difference estimates on (1) Rural versus urban, (2) married women versus married men, (3) before and after 1987. 'Receiving Pension (Y/N)" is an indicator variable equal to one if the individual is receiving pension payments. "Worked (Y/N)" is an indicator variable equal to one if the individual reports working in the reference week. "Hours Worked among Working" is a continuous variable equal to hours worked per week; this sample is limited to individuals who worked in the reference week. "Hours worked per week" is a continuous variable equal to hours worked per week; the sample includes all working and non-working individuals.

the coefficients of interest β_j^{post}, $\forall j \in [1988, 2013]$, measure the average annual treatment effect of the pension expansion on rural labor supply.

In Appendix Figures A.1 to A.3, we test for robustness to high-dimensional fixed effects and various household demographics. First, in Figure A.1 we show robustness to household size and whether the spouse resides in the household. In Figure A.2 we include a polynomial of the individual's age as a control, to account for non-linear trends in age-specific work decisions. Finally, in Figure A.3 we include state-by-year and state-by-rural area fixed effects to account for different trends and policies by state. Our estimates remain robust to the addition of these controls and fixed effects.

Other Pension Reforms. One potential remaining issue would be differential changes in the urban and rural systems: coincident changes in the urban pension system could affect short-term estimates of the labor supply response, while differential reforms of the rural and urban pension system between 1991 and 2013 could affect long-term estimates of the labor supply response. The 1988 constitutional reform and subsequent 1991 Law #8212/9213 made minor coincident reforms to the urban pension scheme that may have influenced the short-term labor supply of urban workers: a requirement for urban workers to quit their current job in order to claim pensions was removed, the minimum benefit increased and working wages taxed, and the work requirement for urban women to gain pension eligibility reduced from 30 to 25 years. The first two of these reforms in the urban scheme bias our short-term difference-in-differences estimates toward zero, but the third potentially increases our estimated coefficients. Any subsequent reform that differentially affected urban and rural areas over the period considered could affect our long-term estimates. Fortunately, the 1988 constitutional reform and subsequent 1991 law actively combined the rural and urban pension systems; as a result, pension reforms or adjustments made after 1991 affected urban and rural pensioners alike. Significant pension reforms that occurred in our period of analysis include a 1998 reform that cut the size of pension benefits received for all pensioners, and a 2003 reform that decreased the generosity of civil servant pensions.

The descriptive analysis in Figure 1 alleviates these concerns regarding differential trends across rural and urban areas due to other pension reforms, since it shows quite starkly that there were no substantial changes in pension provision or labor supply among married women in urban areas, while there were sharp changes to married women's labor supply in rural areas after the pension expansion in 1991. Figure 3 further alleviates the concern that differential trends between rural and urban workers threaten exogeneity by comparing the estimates from specification 1 run on different groups in rural areas – married women and married men. If married women's labor supply in rural areas is primarily influenced by aggregate labor market shocks that were absent from urban areas, the labor supply of married men in rural areas would mimic the patterns found in Figure

2. Figure 3 shows little evidence of similar adjustments in the labor supply of married men immediately after the reform, and some evidence that the decline in rural labor supply in the later years of our sample was a broader rural phenomenon.

Other Nationwide Policies and Shocks. Two additional policies implemented in Brazil during the period of interest that may differentially affect urban and rural areas include Bolsa Familia (BF) and the Benefício de Prestação Continuada (BPC). BF, a large conditional cash transfer to urban mothers who send their children to school, was introduced in 2002 and may have influenced urban mothers' ability or desire to work outside the home.[2] It is not clear in which direction this would affect our estimates: mothers without children to help around the home may have more housework to do, but mothers without children to take care of around the home may have less housework to do. Regardless, the first BF payments were made in 2005 (14 years after the pension reform), and this would only affect the last few years of estimates. Further, the payments primarily affect women with young children who would be in their 20s or 30s – this age group is younger than our cohorts of interest, who were between 25 and 74 in 1991.

The BPC, the second policy of consideration, guaranteed a pension equal to the minimum wage to low-income elderly and disabled individuals. This program, introduced in 1993 and available to rural and urban residents, distributed its first payments in 1996. Low-income individuals were eligible for old-age assistance at age 67 (lowered to 65 in 2003) provided that their family income was no greater than 25 percent of the minimum wage and that they did not receive income from other social security programs or retirement pensions. This implementation of the BPC could affect estimates of the labor-supply response to the 1991 pension expansion if it differentially influenced the labor supply of married rural women. However, Kassouf & de Oliveira (2012) find minimal labor-supply response to the BPC: eligible individuals, age 65 and above, show a small decline in labor supply upon receipt, and do not appear to make anticipatory adjustments in labor supply, while co-residing younger household members do not adjust labor supply when an elderly household member begins to receive the BPC. Our own findings alleviate concerns about endogeneity introduced by the BPC: we find that the largest increases in pension payments and labor supply adjustments were between 1991 and 1993, three years before the first BPC payments were made. Further, the difference-in-discontinuity specification in Section 4 shows a distinct adjustment in labor supply behavior of married women facing the updated pension eligibility age of 55, but *not* not the BPC age-eligibility thresholds of 67 (from 1991 through 2003) or 65 (after 2003).

Finally, the Brazilian hyperinflation of the 1990s could affect estimates if it differentially affected rural and urban areas throughout this period: inflation peaked twice in the 1990s around 6000 percent, once in January of 1990 and a second time in January of

[2]For a detailed overview of Bolsa Familia, see Soares (2011)

1994.[3] Although these inflationary periods, which ended in 1994, do not align perfectly with our policy of interest, they may affect estimates until 1994 if they had different impacts on rural and urban areas. However, according to Baumann (2002), there have been relatively small variations in inequalities between urban and rural areas between 1990 and 1997, which includes the inflationary periods, that alleviate the potential concern.

4.3 Difference-in-discontinuities

We also examine whether introducing a pension eligibility age at 55 influenced the age at which women within a cohort retire. To do so, we use a difference-in-discontinuities specification, and estimate a discontinuity in various outcomes at age 55, over time. Figure 4 shows the discontinuity in pension receipt and labor supply at age 55 among rural women in each year in the three decades surrounding the reform. The first panel shows the stark jump in pension receipt at age 55, after the reform. The subsequent panels look at the probability that women worked, the hours worked among the working, and the hours worked per week.

Prior to the reform in 1991, there was a negligible difference in the probability that a rural woman was working at age 55 versus age 54. However, by 1995, women were four percentage points less likely to work at age 55 than 54, and that difference increased marginally as the years progressed. In Appendix B.2, we test the robustness of this discontinuity by running analogous specifications using ages other than the age of eligibility – 50, 60, and 65 – for men as well as women, and find no evidence of a discontinuous decrease in labor supply at those ages in any year between 1981 and 2006. These findings suggest that individuals who had less than the required years of work experience when the expansion was implemented tried to attain the required years by the age of 55. We capture this empirical reality in our theoretical model by assuming that people within a cohort plan to retire at a certain age, which we call the target retirement age. We then allow the target retirement age to change in response to the pension expansion.

5 Conceptual Framework

This section introduces a conceptual framework in which individuals make lifetime, rather than annual, labor supply decisions to understand why we might see increases in labor supply in response to pension expansion. Individual i from cohort c lives a life of length $\bar{a}_c$ and receives utility from consumption of market-produced goods, C, and of home-produced goods, H, over their lifetime. Their consumption of home goods is inversely proportional to the individual's lifetime supply of market labor, L, with $H = \bar{a}_c - L$. L is the individual's choice variable. Because we are focusing on the labor supply of secondary earners, we assume that the individual holds some household wealth W that

[3] https://www.rateinflation.com/inflation-rate/brazil-historical-inflation-rate/

Figure 4: Difference in Discontinuity at Age 55 among Rural Married Women

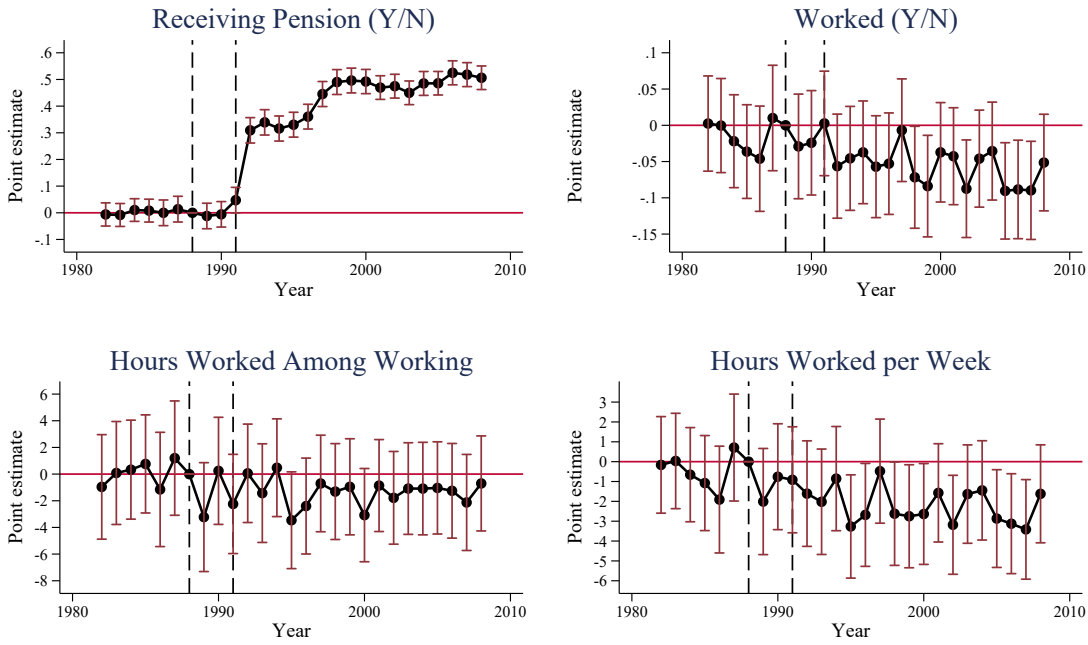

Notes. These graphs show the point estimates and 95% confidence intervals of the difference in discontinuity at age 55, using a bandwidth of 4 years, between 1981 and 2009. The sample is restricted to married women living in rural areas, and the graph shows a discontinuous decrease in employment that develops at age 55 a few years after the reform. 'Receiving Pension (Y/N)" is an indicator variable equal to one if the individual is receiving pension payments. "Worked (Y/N)" is an indicator variable equal to one if the individual reports working in the reference week. "Hours Worked among Working" is a continuous variable equal to hours worked per week; this sample is limited to individuals who worked in the reference week. "Hours worked per week" is a continuous variable equal to hours worked per week; the sample includes all working and non-working individuals.

does not depend on their labor supply. Without the pension regime, individuals receive a market wage, w, that does not change over their lifetime and is not necessarily equal to the worker's marginal product. The pension regime provides the individual with a pension of present discounted value $\tilde{P}$, if they work at least $\tilde{L}_c$ years over their lifetime.[4] People who work when the expanded pension regime is in place receive a market wage w_P, which may or may not be the market wage that prevails without the pension.

The individual's utility of home-produced relative to market-produced consumption, which can be influenced by societal norms and can also be interpreted as the opportunity cost of market work, is captured by α_i and is heterogeneous across individuals. Individuals maximize their utility subject to a lifetime budget constraint:

[4]This life-cycle model abstracts from eligibility age. Under the assumption of zero discounting, people are indifferent with respect to which years in their life they work. Introducing an eligibility age, as we do in section 5.2, encourages them to complete the work requirement before their eligibility age to maximize the present discounted value of their pension benefits.

$$\max_{L} \ (1-\alpha_i) log\ C + \alpha_i log\ (\bar{a}_c - \ L)$$

$$s.t. \quad C \leq \begin{cases} wL + W & \text{without pension} \\ w_p L + W + \left(\tilde{P} \times \mathbb{1}_{L \geq \tilde{L}_c}\right) & \text{with pension} \end{cases}$$

5.1 Lifetime Labor Supply

We find that any pension scheme, $\{\tilde{P}, \tilde{L}_C\}$, affects people's lifetime labor supply differently according to the utility weight they place on home relative to market production, α_i. We illustrate these results with a series of figures that show the number of years an individual chooses to work over their lifetime as a function of the utility they place on home versus market production. Figure 5, Panel A, describes the mechanisms driving the lifetime labor supply response to the pension expansion. First, the increase in the size of the pension creates a wealth effect that encourages all workers, regardless of their preferences, to work fewer years. People with a higher utility weight on home production decrease their market labor supply by more than their counterparts with a lower utility weight on home production as a result of this wealth effect. However, people with a higher utility weight on home production are also more likely to be constrained by the work requirement. This second effect, which we refer to as the eligibility effect, counteracts the wealth effect and encourages some people to work more than they would without the work requirement. Some of those responding to this eligibility effect may even work more than they would have without the pension expansion. When enough people in the population experience an eligibility effect that dominates the wealth effect, the pension expansion increases aggregate labor market participation.

These dynamics create four groups of workers, identified in Panel B of Figure 5, that respond differently to the pension expansion. People with a low value of home production who worked prior to the pension expansion, $\alpha_i \leq \alpha_1 = \frac{w_p(\bar{a}_c + W - \tilde{L}_c)}{w_P(\bar{a}_c + W) + \tilde{P}}$, reduce their labor supply when the pension is available, but continue to work more than the minimum number of years required to achieve pension eligibility. These individuals, who we refer to as "market workers," respond to the wealth effect created by the additional pension wealth and are not constrained by the pension's minimum work requirement. People with a slightly higher value of home production, $\alpha_i \in \left(\alpha_1, \alpha_2 = \frac{\bar{a}_c + W - \tilde{L}_c}{\bar{a}_c + W}\right)$ similarly reduce their labor supply when the pension is available, but are constrained by the minimum work requirement and thus work exactly $\tilde{L}_c$ years over their lifetime. These individuals, who we call "down-compliers," respond to the minimum work requirement as well as the wealth effect; they work less than they would have without the pension but more than they would have in the absence of the work requirement. A second group of compliers, "up-

Figure 5: Heterogeneous Lifetime Labor Supply Responses to Pension Expansion

A. Wealth and Eligibility Effects

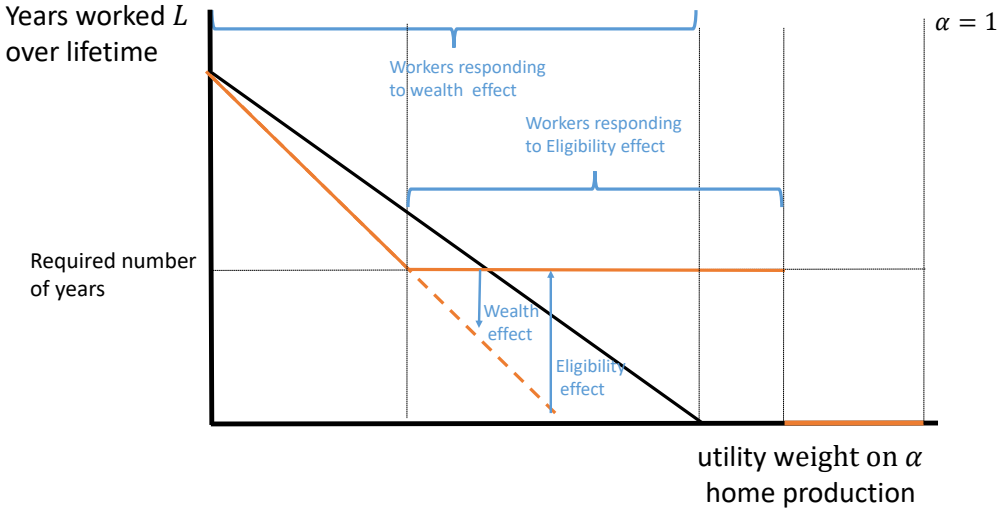

B. Categories of Worker Responses

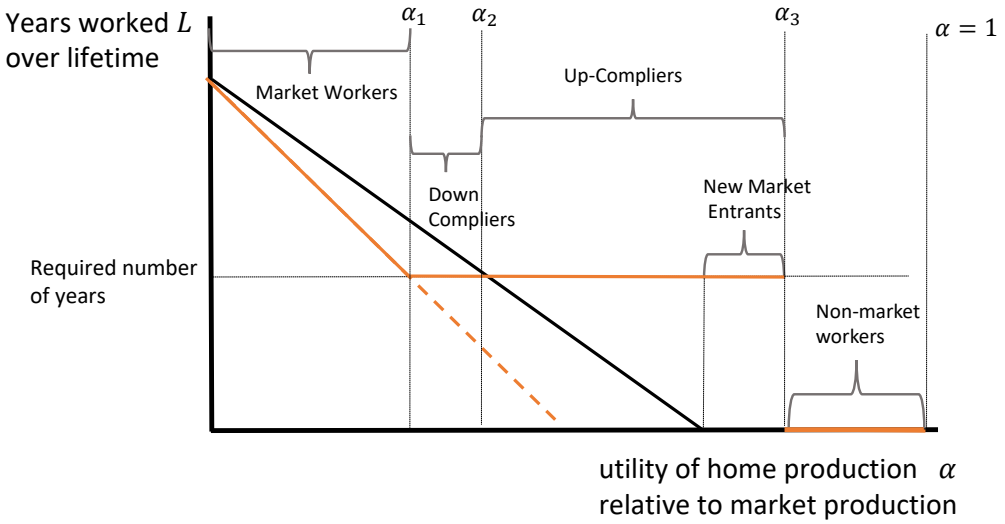

Notes. This figure shows the optimal number of years an individual works over their lifetime, for heterogeneous utility weight on home production ranging from $\alpha_i \in (0, 1)$. The black solid line describes optimal years worked without a pension expansion. The orange dotted line describes optimal years worked under a pension expansion with no work requirement. The solid orange line describes optimal years worked under a pension expansion with a work requirement. Panel A describes both the traditional wealth effect arising from more generous pension benefits, decreasing lifetime labor supply, and the eligibility effect that increases lifetime labor supply. Panel B identifies four categories of worker responses. Market workers respond only to the wealth effect, and thus decrease their lifetime labor supply following the reform. Down-compliers respond to both the wealth effect and the eligibility effect; the wealth effect dominates and therefore down-compliers decrease their lifetime labor supply. Up-compliers similarly respond to both effects, but the eligibility effect dominates, encouraging up-compliers to increase their lifetime labor supply. Non-market workers do not change their labor supply as a result of the pension expansion.

compliers" with a slightly higher $\alpha_i \in (\alpha_2, \alpha_3)$, also works exactly $\tilde{L}_c$ years.[5] Up-compliers similarly respond to both the minimum-work requirement and the wealth effect, but work more under the pension regime than they would have without the pension. This group includes new market entrants who would not have worked without the pension, as well as previous market workers who increase the number of years they work. Finally, individuals with a high value of home production, $\alpha_i > \alpha_3$, who we call "non-responders," do not adjust the number of years they plan to work under the pension expansion. Through these mechanisms, this pension expansion encourages people who are more attached to the labor force to decrease their lifetime labor supply, and those who are less attached to the labor force to increase their lifetime labor supply.

5.2 Annual Labor Supply

Much empirical work on pension expansions, including our own, estimates their impact on annual, rather than lifetime, labor supply. To adjust our model to accommodate annual treatment effects, we add three assumptions. First, we assume that people are eligible to receive the rural pension, provided that they have worked at least $\tilde{L}_c$ years, only when they reach eligibility age, $\tilde{a}_E$. The pension scheme is now described by the triple $\mathcal{P} = \{\tilde{L}, \tilde{P}, \tilde{a}_E\}$. Second, we assume that each cohort has a target retirement age – ie., an age at which all members of the cohort plan to and will retire – prior to the pension expansion, $\bar{a}_R$, that could be different from its target retirement age following the expansion, $\tilde{a}_R$.[6] Finally, we assume that $\alpha_i \sim G(\alpha)$ within a cohort, and that aggregate lifetime labor supply within the cohort is smoothed across the years before the cohort reaches its target retirement age. As a result, if a pension expansion is introduced in year j when individuals from cohort c are age a_{cj}, the remaining labor supply for that cohort is equally distributed across the next $\tilde{a}_R - a_{cj}$ years.

Decomposing the lifetime labor supply response in this way introduces a third channel, in addition to the wealth and eligibility channels, by which the pension expansion can influence labor supply at an annual level. That channel, which we call the retirement-timing channel, arises because the expansion may encourage people to re-allocate their labor supply over their lifetime. The following equation, derived in detail in appendix A, describes the average annual treatment effect for cohort c:

[5] The solution to α_3 is detailed in Appendix A.1

[6] This target retirement age is consistent with literature showing reference-dependents retirement age and our lifetime model. Further, our assumption of zero discounting on both home and market consumption, will encourage people to concentrate their working years earlier in their career so that they are eligible for the pension as soon as possible.

$$\Delta L_{ct} = \frac{1}{\tilde{a}_R - a_{cj}} \Bigg[\int_{\alpha} \underbrace{(1-\alpha_i)(\bar{a}_c + W)\left(\frac{\bar{a}_R - \tilde{a_R}}{\bar{a}_R}\right)}_{\text{Retirement Timing Effect}} \underbrace{-\frac{\alpha_i \tilde{P}}{w_P}}_{\text{Wealth Effect}} \, dG(\alpha) + \int_{\alpha_1}^{\alpha_3} \underbrace{\tilde{L}_c - (1-\alpha_i)(\bar{a}_c + W) + \frac{\alpha_i \tilde{P}}{w_P}}_{\text{Eligibility Effect}} dG(\alpha) \Bigg] \tag{2}$$

This expression describes the three mechanisms through which the pension expansion influences average cohort-level annual labor supply: the negative wealth effect, the positive eligibility effect, and the retirement timing effect – which could be positive or negative in any given year but aggregates to zero over the lifetime. The distribution of α_i within a cohort determines which effect dominates in any given year. While the lifetime model of labor supply presented in Section 5.1 describes the optimal amount of work someone would choose if they knew, at the beginning of their career, the pension scheme that would prevail when they were ready to retire, this decomposition captures the annual adjustments workers make, at various points in their working life, when a new pension regime is unexpectedly introduced in the middle of their career.

5.3 Bringing the Model to the Data

The coefficient estimates, β_j^{post} for $j > 1991$, found by running specification 1 on a sample restricted to individuals in cohort c who were of working age in 1991, estimate the theoretical average annual treatment effect, ΔL_{ct}, described in Equation 2. Table 2 groups observations into five-year bins to present the extensive-margin empirical treatment effect for five different cohorts, defined by their age upon the implementation of the pension expansion in 1991. The years between the 1988 Constitution and the 1991 implementation of the pension expansion serve as the reference period. In this table, extensive-margin estimates of ΔL_{ct} range between seven to nine percentage points for younger cohorts, and rise to 12 to 19 percentage points for older cohorts. These annual adjustments include the eligibility effect and reallocation of labor supply throughout the working-life due to adjustments in retirement age, as well as the standard wealth effect.

Building on the theoretical model, we then calculate the realized change in lifetime work for each cohort, ΔL_c, by aggregating annual treatment effects for all married women younger than 70 within a cohort, over the available time period: $\Delta L_c = \sum_{t=1991}^{2012} N_{ct} \Delta L_{ct}$. The direction of this estimate indicates whether the negative wealth effect on market workers and down-compliers dominates the positive eligibility effect on up-compliers within a cohort. The final rows of Table 2 present the empirical analogue of ΔL_c for the five cohorts of interest.[7] The increase in lifetime work realized by 2012 suggests that

[7] See Appendix C for details on the calculation of years per worker and percent increase.

the eligibility effect dominates the wealth effect for all cohorts.

Table 2: Treatment effect on extensive margin among married women by cohort

	(1) 25-34 in 1991	(2) 35-44 in 1991	(3) 45-55 in 1991	(4) 55-59 in 1991	(5) 60-69 in 1991
Rural X 1982-1986	0.01	0.00	-0.01	-0.01	0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Rural X 1992-1996	0.07**	0.09**	0.15**	0.17**	0.16**
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
Rural X 1997-2001	0.07**	0.11**	0.18**	0.15**	0.12**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)
Rural X 2002-2006	0.09**	0.16**	0.19**	0.14**	
	(0.02)	(0.02)	(0.02)	(0.03)	
Rural X 2007-2012	0.08**	0.14**	0.17**		
	(0.02)	(0.02)	(0.02)		
Observations	395321	373183	214262	60476	58244
R^2	0.034	0.020	0.036	0.034	0.026
Years of Work Required	15	10-15	5-10	5	5
Age in 2012	46-55	56-65	66-75	76-80	81-90
Realized Lifetime adjustments					
Realized years per worker	3.12	4.07	4.89	3.09	1.94
Pct Increase in worker-years	6.52	8.14	10.49	7.64	4.80
Years per person	1.44	2.15	2.53	1.40	0.72
Potential Lifetime adjustments (years per person)					
Upper bound, $\tilde{a}_R = 70$	3.64	4.81	3.73	2.02	0.87
Estimate with $\tilde{a}_R$ set at discontinuity	2.30	2.48	1.43	0.85	0.80

Notes. Standard errors, clustered by rural and state, in parentheses. * $p < .10$, ** $p < .05$. This table shows the β_j coefficients estimated from the regression:
$y_{it} = \alpha \text{Rural}_r + \sum_{j=1}^{5} \beta_j \text{YearCat}_j \times \text{Rural}_r + \delta_t + \mu_s$, where YearCat_j groups observations into five-year bins, with the years 1988-1991 serving as the reference period. δ_t and μ_t represent time and state fixed effects. The regression is run for five different cohorts, defined by their age upon implementation of the pension expansion in 1991. The realized change in lifetime work is calculated by aggregating the treatment effects across the population from 1992 through 2012. Per worker and per person estimates are calculated by dividing aggregate estimates by the number of workers or size of population within cohort in 1991. Potential lifetime adjustments are calculated as described in Section 6.1.

The conceptual framework also captures the average annual treatment effect in any given year, ΔL_t, estimated empirically by restricting the sample underlying specification 1 to all working-age individuals as in Figure 2, Panel A. The resulting difference-in-differences coefficients on labor supply are not cohort-specific but vary over time. The analog in the model is $\Delta L_t = \sum_c N_c \Delta L_{ct}$, where N_c is the population share in cohort c. The theoretical interpretations of the extensive-margin and aggregate treatment effects suggest a strong eligibility effect that dominates the wealth effect at the population level, which we explore further in section 6. To better understand the intensive-margin decline in hours worked among the working, we tweak the model to include a discrete working choice (full, part-time, or none) in Appendix A, and explore empirical changes in part-time labor supply in section 7.

6 Changes in Labor Market Participation by Cohort

The large increases in married women's overall labor supply at the annual frequency, shown in Figure 2, were driven by large, steady increases in labor market participation among the cohorts aged 25 to 69 when the pension expansion was implemented in 1991. Table 2 summarizes these annual labor market participation patterns by five-year increments for the five cohorts of interest. While these observed annual increases in labor market participation are larger among older cohorts, they do not allow us to conclude that older married women increased their lifetime labor supply participation by more than their younger counterparts. Younger cohorts knew the parameters of the pension regime that would prevail upon their retirement at the beginning of their careers, but older cohorts had to adjust to a new pension regime in the middle of their careers. As a result, older cohorts had fewer years to accumulate additional years of work before their target retirement age, contributing to annual labor supply adjustments that are larger in magnitude for those closer to retirement eligibility age than those of younger cohorts.[8]

6.1 Bounding Lifetime Labor Supply Adjustments at the Cohort Level

Fully understanding the impact of the pension expansion on married women's labor market participation requires aggregating cohorts' annual responses into lifetime measures. The final rows of Table 2 present realized lifetime labor supply – that is, the average increase in years worked over the lifetime among workers in each cohort from 1991 through 2012. Women who were between the ages of 25 and 34 in 1991, for example, increased their lifetime labor supply by an average of 1.44 years per person by 2012, while women who were between the ages of 45 and 55 increased their lifetime labor supply by an average of 2.53 years per person by 2012. Realized lifetime increases are then noticeably smaller for cohorts who were older than eligibility age upon implementation, falling to an average of 1.4 years per person for those between 55 and 59 in 1991, and of 0.72 years per person for those between 60 and 69 in 1991.

These realized lifetime increases are complete for older cohorts that have fully retired, but younger cohorts may not have completed their working lives by 2012. We next conduct a bounding exercise, motivated by two implications of Equation 2, to provide a range for the lifetime adjustments in labor supply for each cohort. First, all annual adjustments for which the individual is older than the target retirement age, $a_{cj} > \tilde{a}_R$, are weakly negative among cohorts with a dominant eligibility effect. Second, annual labor supply adjustments are similar in magnitude prior to the adjusted target retirement age.[9]

[8] This theoretical prediction manifests in Equation 2, in which the annual labor supply response is scaled by $\frac{1}{\tilde{a}_R - a_{cj}}$, where a_{cj} is the age of cohort c upon introduction of the pension expansion and $\tilde{a}_R$ is the cohort's retirement age.

[9] People are re-allocating their labor from later in life to before their new target retirement age. We

An upper bound on lifetime labor supply is then the sum of all annual increases in labor supply, realized or projected at a steady magnitude, prior to the cohort's adjusted target retirement age. Suppose, for example, that all cohorts adjust their target retirement age to 70, which is the required retirement age for urban workers, following the expansion. The penultimate row of Table 2 presents the implied upper bounds on lifetime adjustments for each cohort by aggregating the annual estimates of labor market participation adjustments from their age at implementation to age 70.

Alternatively, the difference-in-discontinuities exercise in Section 4.3 can inform assumptions on cohorts' adjusted retirement age. Echoing the discontinuity in labor force participation that developed among cohorts turning 55 in 1998 and later, we assume that the target retirement age ($\tilde{a}_R$ in our model) fell to 55 for cohorts that were 48 and younger in 1991. The target retirement age of older cohorts, however, was potentially limited by the design of the work requirement: women that were 55 in 1991, for example, were required to work for at least 5 years to achieve pension eligibility. Any 'new market entrant,' as Figure 5 defines, in that cohort would not be able to retire with pension access until age 60. For these older cohorts, we assume a target retirement age equal to the youngest possible age at which an individual in the cohort could retire and receive the lifetime pension, had they not worked at all prior to the pension expansion.[10] Finally, we assume steady annual increases among each cohort between pension implementation and target retirement age, equal to the highest cohort-level annual increase in Table 2. The final rows of that table present the estimates of potential lifetime adjustments calculated under these assumptions for five cohorts of interest.

The increase in labor market participation for cohorts aged 25-34 and 35-44 in 1991, for example, is bounded above by 3.64 and 4.81 years per person, respectively, if all workers within the cohort work until age 70. However, if workers in those cohorts retire as soon as they are age-eligible, the lifetime increase in labor market participation is projected to be 2.30 and 2.48 years per person. The increase in labor market participation for the cohort just below eligibility age upon pension implementation, age 45-55 in 1991, is bounded above by 3.73 years per person but estimated at 1.43 years per person when each member of the cohort retires as soon as a new market entrant would become age-eligible. We also calculate the upper bounds on the lifetime increase in labor market participation for cohorts who were older than eligibility age in 2012. The upper bounds for married women in these cohorts are 2.02 years per worker for women who were 55-59 in 1991, and .87 years per person for women who were 60-69 in 1991.

provide alternate bounding exercises based on different assumptions on the target retirement age in Appendix C.

[10]See Appendix C for details on the graduated work requirement, assumed target retirement ages by cohort, and a detailed description of the bounding exercise.

6.2 What Theoretical Mechanisms Could be Driving These Differences in Response across Cohorts?

The observed pattern of lifetime increases in labor force participation across the three youngest cohorts suggests that young women of working age upon implementation of the expansion increased their lifetime labor supply participation by less than their older counterparts, who were also of working age upon implementation, but nearer to the new pension eligibility age. The conceptual framework presented in Section 5.1 suggests two potential mechanisms underlying this empirical pattern. First, the distribution of α among younger cohorts may have been skewed further to the left than that of older cohorts. This would be consistent with changing societal norms or preferences by which younger women place less value in home-produced relative to market-produced goods. In this case, there would be fewer new market entrants and up-compliers (people with a dominant eligibility effect), and more down-compliers and market workers (people with a dominant wealth effect), among younger cohorts than among their older counterparts.

Alternatively, the graduated work requirement – a key aspect of the policy design – could be driving these differences in lifetime labor market participation responses across cohorts of working age upon pension implementation, even when the distribution of α is constant across cohorts. The 15-year work requirement, referred to as $\tilde{L}_c$ in Section 5, was phased in gradually across cohorts; women who were age-eligible in 1991 were required to produce evidence of only five years of rural work, while those who became age-eligible in 2011 and later were required to produce evidence of 15 years of work to access the pension. Our model predicts that a shorter work requirement creates more lifetime years of work among populations with high home productivity, because it brings more people into the labor force – even if some in the labor force work for fewer years than they would have under a longer work requirement. Figure 6 illustrates this theoretical result by comparing the impact on lifetime labor supply of the pension expansion with a 15-year work requirement (Panel A) to that of a pension expansion with a five-year work requirement (Panel B), analogous to comparing the impact of the policy parameters facing the younger cohorts versus those facing the older cohorts. The five-year work requirement has a dominant negative wealth effect for a broader range of home productivity than the 15-year work requirement ($\alpha_2' > \alpha_2$). However, the five-year work requirement also brings more people with higher levels of home productivity into the labor force than does the 15-year requirement ($\alpha_3' > \alpha_3$). If cultural or norm-based preferences meant that the distribution of home productivity among married women in rural Brazil was skewed right, with high density between α_3 and α_3' in Figure 6, the design of the 1991 phase-in could have led to larger increases in lifetime labor supply among older cohorts with lower work requirements than their younger counterparts with higher work requirements.

Though the lack of reliable data that tracks individuals' informal work history in rural

Figure 6: Heterogeneous Lifetime Labor Supply Responses to Pension Expansion with Different Work Requirements

A. Worker Responses to 15-year Work Requirement

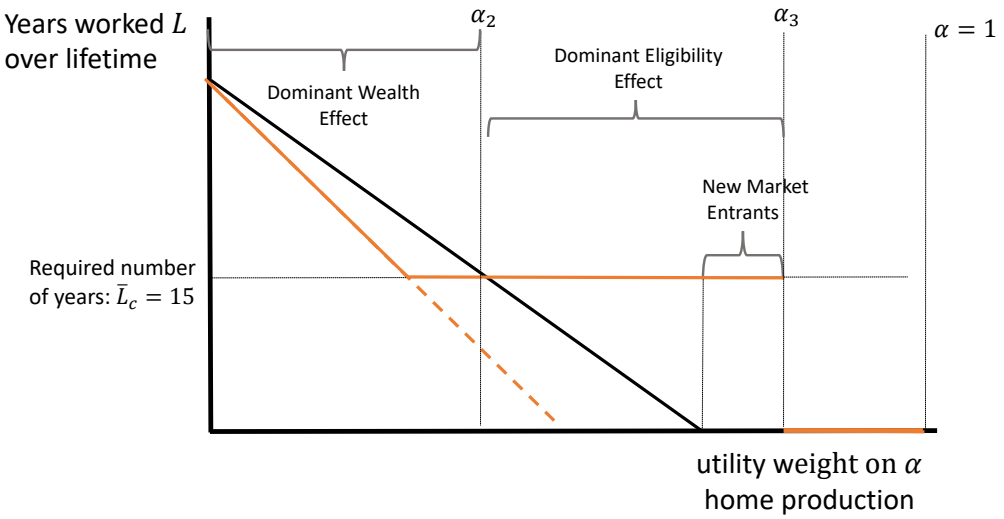

B. Worker Responses to 5-year work requirement

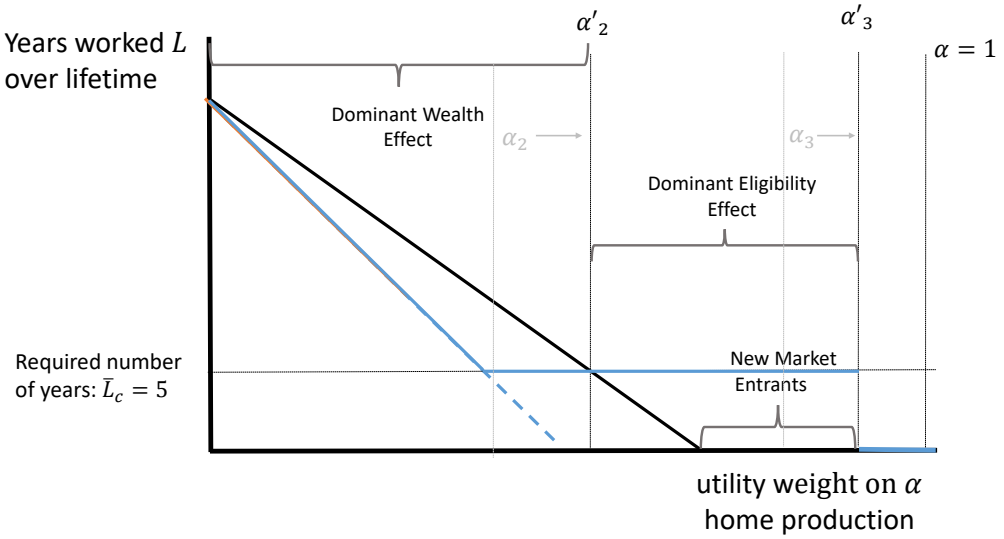

Notes. This figure shows the optimal number of years worked over an individual's lifetime, for heterogeneous home productivity ranging from $\alpha_i \in (0, 1)$. The black line shows optimal years worked without a pension expansion. The dotted lines show optimal years worked under an expansion with no work requirement. The solid orange and blue lines describe optimal years worked under expansions with various work requirements. Figure A describes the labor supply response to a pension expansion with a 15-year work requirement. People with $\alpha_i < \alpha_2$ experience a dominant wealth effect and decrease their lifetime labor supply as a result of the expansion. People with $\alpha_i \in (\alpha_2, \alpha_3)$ experience a dominant eligibility effect and increase their lifetime labor supply as a result of the expansion. Figure B describes the labor supply response to a pension with a five-year work requirement. Under a five-year work requirement, the wealth effect dominates for a wider range of α, with $\alpha'_2 > \alpha_2$, than under a 15-year work requirement. However, the eligibility effect induced by a 5-year work requirement also brings more people into the labor force than does the 15-year work requirement, with $\alpha'_3 > \alpha_3$.

Brazil prevents us from precisely estimating the distribution of α among rural women, the available data indicate a distribution that is skewed right. We find little evidence of a shift in preference distribution when we calculate changes in lifetime labor market participation among younger cohorts with the same work requirement.[11] Large increases in extensive-margin lifetime labor supply among all cohorts imply that the increase in labor supply among up-compliers is larger than the decrease in labor supply among down-compliers and market-workers. Many married women are at the cusp of labor market participation prior to pension expansion, and are willing to spend a small number of years in the labor force in exchange for a lifetime pension.

7 Part-time Work and the Intensive Margin

We next consider intensive-margin adjustments in labor supply that occurred following the reform. Empirical estimates presented in the third graph of each panel in Figures 2 and 3 show that the average length of the workweek fell dramatically among all cohorts of rural women. Figure 7 explores this result in detail. In Panel A, the fraction of married women participating in the market labor force who worked part-time (less than 30 hours per week) increased dramatically in rural areas relative to urban areas between 1990 and 1992. This increase occurred among all working-age cohorts in 1991, and part-time work remained high throughout the 2000s.[12] Panel B describes the broader distribution of working hours among all married women in rural areas. The fraction of married women in rural areas who were working full-time, between 30 and 44 or over 45 hours per week, remained steady throughout the period shown. At the same time, the fraction of married women in rural areas working between one and 29 hours increased by more than ten percentage points between 1990 and 1992. Eighty-one percent of this increase was due to an increase in married women working very few hours, between one and 19, each week. These patterns suggest that the large increase in part-time work, and therefore the decline in average length of the workday, was driven by new market entrants working part-time, rather than by existing market workers decreasing their working hours.

Extending our model to accommodate part-time work sheds light on these adjustments at finer hourly increments.[13] Women working full-time in the absence of the pension expansion would be market workers or down-compliers. Thus, women who cut back

[11] Appendix C calculates lifetime labor supply in smaller increments among cohorts with the same work requirement, and finds similar adjustments across these cohorts. This suggests that it is likely the difference in work requirement, rather than in preferences, across cohorts that drives cohort-level participation patterns.

[12] Appendix D.2 explores this result more structurally, with a difference-in-difference estimation. Part-time work among working married women in rural areas increased by 14-19 percentage points relative to their urban counterparts following the pension expansion.

[13] The discrete part-time model in Appendix A.3 predicts that compliers are more likely to choose to work part-time. Alternatively, interpreting L as the number of hours worked per year indicates that down-compliers will work longer (potentially part-time) days than up-compliers.

Figure 7: Part-time Work Among Married Women

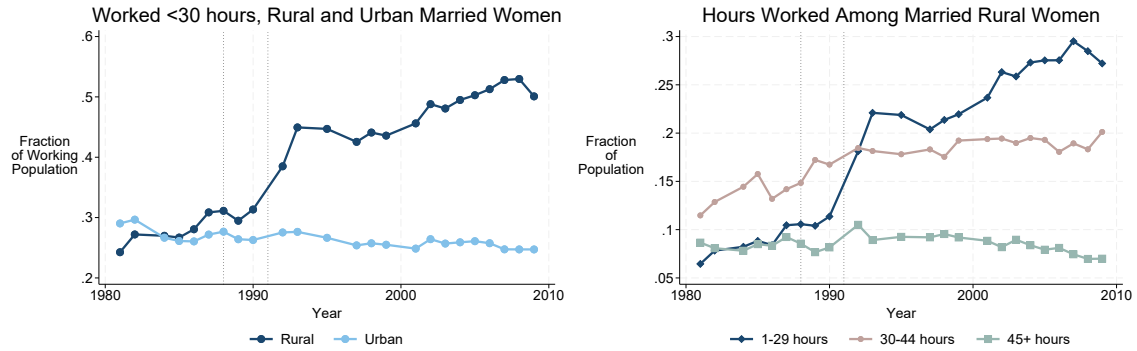

Notes. The figure describes trends in part-time work among married women between 25-69. Panel A shows the percent of the population that reported working in the reference week who worked less than 30 hours per week in rural (dark lines) and urban (light lines) areas. Panel B shows the percent of the rural married female population that reported working in the reference week who worked between 1 and 29 hours, 30 and 44 hours, and above 44 hours per week. Vertical dashed lines indicate the years 1988, when the constitution was passed, and 1991, when the pension expansion was implemented.

from 45 weekly hours, or from between 30 and 44 weekly hours, are market workers and down-compliers responding to a dominant wealth effect. Women not working at all, or working very few hours in the absence of the pension expansion, are likely to be up-compliers or non-respondents. The increase in people working between 1 and 19 hours, is then likely driven by new market entrants, or "up-compliers," responding to a dominant eligibility effect. Like our extensive-margin results, these intensive-margin results suggest a distribution of home productivity among married women, or α, that is skewed right.

8 Gender Differences in Labor Supply Response

While the focus has been on married women, one might also expect such a large expansion in pension benefits to influence the labor market behavior of married men. The expansion was not explicitly limited by gender, yet the economic and cultural context in which the reform was implemented implied that it had substantially different impacts on the labor supply of men and women. Prior to the 1991 reform, 90% of rural households in Brazil reported a man as their household head, and over 90% of married rural men participated in the market labor force. In contrast, only 25% of married rural women participated in the market labor force (Figure 3). Households could receive only one rural pension prior to the reform, and age-eligible men were the primary recipients of that pension. The 1991 pension reform made rural pensions more generous for men by doubling the annual payout and decreasing the age at which men could start receiving benefits, to 60 from 65. The work requirement also increased for men, as it did for women, but most men were already working more years than required under the new rules. Accordingly, while both occurred for men, the increase in pension generosity was not as large for men as it was

Figure 8: Adjustments in the Workweek Among Women and Men

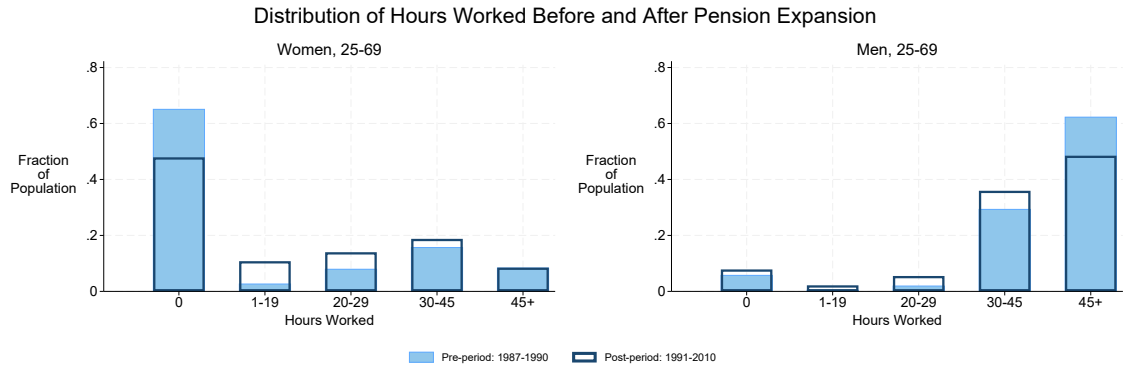

Notes. This figure shows the distribution of full- and part-time work among women and men of working age before and after implementation of the expansion. Both figures separate the population age 25-69 into one of four categories based on hours worked per week: not working (0 hours), low-hour part time work (1-19 hours), high-hour part-time work (20-29 hours), full-time work (40-44 hours), and high-hour full dime work (45+).

for women and the increase in the work requirement was not as salient.

That said, these reforms were sizeable enough to influence men's labor market behavior. The difference-in-difference estimates in Figure 3, Panel B, reveal a slight increase in annual labor force participation among married men following the reform. This increase, around 1 percentage point annually, is substantially smaller than that of women and falls to zero within a few years. The third graph in each panel of Figure 3 shows that hours worked among the working decreases for both rural men and women, relative to their urban counterparts. The difference-in-differences estimates confirm that this decline is substantially smaller among men than women. The final graph of each Panel in Figure 3 shows that, while the long-term annual adjustments in total hours worked among the full populations are quite similar, men – unlike women – do not show a clear, immediate increase in overall hours worked upon pension implementation. These findings are consistent with de Carvalho Filho (2008), who studies the impact of this pension expansion on the labor supply behavior of newly age-eligible men. She shows that newly age-eligible men, 60-64 in 1991, decreased their labor market participation by 38%, and reduced total hours worked by 22.5, between 1990 and 1992. Her focus on the short-term response of newly age-eligible men abstracts from the broader lifetime adjustments in labor supply treated within this paper, and explains why she finds negative labor-market adjustments in the short-term that are so much larger than those we present in Section 4.2.1.[14]

Figure 8 divides the working-age population of each gender by non-market participants, part-time workers, and full-time workers. The two groups of workers are further

[14]We discuss these differences in results across the two papers in detail in Appendix D.1. The key differences are that the de Carvalho Filho (2008) paper focuses on short-term adjustments, and her choice to use different age-groups of rural males as an additional comparison groups for a third difference.

separated into part-time workers who work between one and 19 hours versus those who work between 20 and 29 hours per week; and full-time workers who work between 30 and 44 hours versus those who work 45 or more hours per week. This figure provides suggestive evidence on the distribution of utility weight men place on home relative to market production, in comparison to that of women, and the dominant "types" of worker responses to the pension expansion across genders. The adjustment among men shows strong evidence of left-skewed distribution and a negative wealth effect: substantially fewer men work more than 45 hours a week following the reform, with that decline corresponding to increases throughout the lower end of the hours distribution. Women, on the other hand, show evidence of a right-skewed distribution and strong eligibility effect: more women enter the labor force following the expansion, and those entrants are concentrated among part-time workers – and among low-hour part-time workers in particular. The fact that 82 percent of the increase in part-time work among women was due to an increase in those working 1-19 hours (versus 20-29), suggests that many of these workers are up-compliers who are incentivized to do the minimum amount of work required to earn a pension, but in fact place a high weight on home relative to market production.

9 Conclusion

Middle-income countries, like Brazil, often have low female labor force participation. Cultural, norm-based, and market-based frictions that restrict access to the economy for half the population hinder unleashing its potential for growth and future development. Policies that, with or without intent, expand female labor force participation at such an enormous rate are difficult to find. We document a substantial shift in women's labor-market participation in response to the pension scheme. Our results suggest that the pension reform played an important role in equalizing market forces for men and women.

This paper sheds light on the willingness and the ability of workers to react to retirement incentives in a forward-looking manner. The results regarding the immediately eligible cohort indicate that elderly workers have the ability to increase their labor supply given the right incentives, and the results regarding the younger cohorts indicate that retirement policies enacted today may have unforeseen effects among those who are not currently eligible for benefits, but will be in the future. While the responses seem to be much larger than what has been documented in other contexts, the pensions were indeed sizeable. Given how important the equivalent of an additional minimum wage might have been to a low-income rural family, it may be more reasonable to expect that the work requirements of the pension could dominate the wealth effect on labor supply.

Our study looks at the impact of the pension expansion and coincident rural work requirements on work by married women living in rural areas. It would be interesting to further explore the interplay between working outside the household, marriage decisions, and time spent with children. For instance, whether households were more likely to hire

domestic workers as women did less domestic work to engage in paid rural work outside of the house. There is evidence that an expansion of rural pensions in other countries influence, for example, the labor supply decisions of adult children or the education support of younger children. We would like to explore these impacts in the case of Brazil.

This study adds more broadly to the literature regarding labor supply responses to retirement policies. Reforms of benefits and social security often demand an analysis of the associated labor supply responses among the eligible cohort. However, this paper shows that an expansion of benefits can, under some circumstances, *increase* labor supply if qualifications are properly managed.

References

Anderson, S. & Eswaran, M. (2008). What Determines Female Autonomy? Evidence from Bangladesh. *Journal of Development Economics*, 90(2), 179–91.

Bando, R., Galiani, S., & Gertler, P. (2016). Non-contributory Pensions. *Labour Economics*, 38, 47–58.

Bando, R., Galiani, S., & Gertler, P. (2020). The Effects of Noncontributory Pensions on Material and Subjective Well-being. *Economic Development and Cultural Change*, 68(4), 1233–1255.

Bando, R., Galiani, S., & Gertler, P. (2022). Another Brick on the Wall: On the Effects of Non-contributory Pensions on Material and Subjective Well Being. *Journal of Economic Behavior & Organization*, 195, 16–26.

Bastian, J. (2020). The Effects of the EITC on Maternal Labor Supply.

Baumann, R. (2002). Brazil in the 1990s: an Economy in Transition. *En: Brazil in the 1990s: an Economy in Transition-London: Palgrave/St. Antony's College, 2002-p. 1-38*.

Behaghel, L. & Blau, D. M. (2012). Framing Social Security Reform: Behavioral Responses to Changes in the Full Retirement Age. *American Economic Journal: Economic Policy*, 4(4), 41–67.

Blau, F. & Kahn, L. (2013). Female Labor Supply: Why is the US Falling Behind? *AEA: Papers and Proceedings*, 103(3), 251–256.

Brown, K. M. (2013). The Link between Pensions and Retirement Timing: Lessons from California Teachers. *Journal of Public Economics*, 98, 1–14.

Coile, C. & Gruber, J. (2007). Future Social Security Entitlements and the Retirement Decision. *The Review of Economics and Statistics*, 89(2), 234–246.

Danzer, A. M. & Zyska, L. (2022). Pensions and Fertility: Micro-Economic Evidence. *American Economic Journal: Economic Policy*.

de Carvalho Filho, I. E. (2008). Old-age Benefits and Retirement Decisions of Rural Elderly in Brazil. *Journal of Development Economics*, 86(1), 129–146.

Dinkelman, T. & Ngai, R. (2021). Home Production, Women's Market Work, and Structural Transformation. *STEG Pathfinding Paper*.

Duque, D. (2021). Rural Pensions and Household Welfare in Brazil. *Working Paper*, (pp. 1–14).

Goldin, C. (1995). The U-Shaped Female Labor Force Function in Economic Development and Economic History. *Investments in Women's Human Capital*, (pp. 61–90).

Gruber, J., Kanninen, O., & Ravaska, T. (2022). Relabeling, retirement and regret. *Journal of Public Economics*, 211, 104677.

Huang, W. & Zhang, C. (2021). The Power of Social Pensions: Evidence from China's New Rural Pension Scheme. *American Economic Journal: Applied Economics*, 13(2), 179–205.

Kassouf, A. L. & de Oliveira, P. (2012). Impact evaluation of the Brazilian non-contributory pension program Benefício de Prestação Continuada (BPC) on family welfare. *Partnership for Economic Policy working paper*, (2012-12).

Kaushal, N. (2014). How Public Pension Affects Elderly Labor Supply and Well-being: Evidence from India. *World Development*, 56, 214–225.

Kleven, H. (2019). *The EITC and the Extensive Margin: A Reappraisal*. Technical report, National Bureau of Economic Research.

Lalive, R., Magesan, A., & Staubli, S. (2023). How social security reform affects retirement and pension claiming. *American Economic Journal: Economic Policy*, 15(3), 115–150.

Manoli, D. S. & Weber, A. (2016). *The effects of the early retirement age on retirement decisions*. Technical report, National Bureau of Economic Research.

Mastrobuoni, G. (2009). Labor Supply Effects of the Recent Social Security Benefit Cuts: Empirical Estimates using Cohort Discontinuities. *Journal of Public Economics*, 93(11-12), 1224–1233.

Neumark, D. & Song, J. (2013). Do Stronger Age Discrimination Laws Make Social Security Reforms More Effective? *Journal of Public Economics*, 108, 1–16.

Samwick, A. A. (1998). New Evidence on Pensions, Social Security, and the Timing of Retirement. *Journal of Public Economics*, 70(2), 207–236.

Seibold, A. (2021). Reference points for retirement behavior: Evidence from german pension discontinuities. *American Economic Review*, 111(4), 1126–1165.

Shu, L. (2018). The Effect of the New Rural Social Pension Insurance Program on the Retirement and Labor Supply Decision in China. *The Journal of the Economics of Ageing*, 12, 135–150.

Soares, F. V. (2011). Brazil's Bolsa Família: a Review. *Economic and Political Weekly*, (pp. 55–60).

Staubli, S. & Zweimüller, J. (2013). Does Raising the Early Retirement Age Increase Employment of Older Workers? *Journal of Public Economics*, 108, 17–32.

Stock, J. H. & Wise, D. A. (1990). Pensions, the Option Value of Work, and Retirement. *Econometrica*, 58(5), 1151–1180.

Verho, J., Hämäläinen, K., & Kanninen, O. (2022). Removing Welfare Traps: Employment Responses in the Finnish Basic Income Experiment. *American Economic Journal: Economic Policy*, 14(1), 501–22.

Whitmore Schanzenbach, D. & Strain, M. R. (2021). Employment Effects of the Earned Income Tax Credit: Taking the Long View. *Tax Policy and the Economy*, 35(1), 87–129.

Appendix

Table A.1: The Brazilian Social Security System for Men and Women, before and after reform

Women	Rural	Urban
Pre Reform		
Eligibility status	* Women who are household heads * Age 65 and above	* All women * Age 70 and above, or 30 years of work
Benefit amount	* 50% of minimum wage	* Min. benefit: 90% of minimum wage * Rises with work and earnings
Work requirement	* Must document rural work for at least 1 of past 3 years	* Must quit job to receive benefits
Contribution requirement	* None	* None
Post Reform		
Eligibility status	* All women * Age 55 and above	* All women * Age 70 and above, or 25 years of work
Benefit amount	* 100% of minimum wage	* Min. benefit: 100% of minimum wage * Rises with work and earnings
Work requirement	* Min. work requirements increase from 5 years in 1991 to 15 years in 2011	* No requirement to quit current job to receive benefits
Contribution requirement	None	* Working wages are taxed

Men	Rural	Urban
Pre-reform		
Eligibility status	* Men who are household heads * Age 65 and above	* All men, regardless of household status * Age 70 and above, or 30 years of work
Benefit amount	* 50% of minimum wage	* Min. benefit: 90% of minimum wage * Rises with work and earnings
Work requirement	* Must document rural work for 1 of past 3 years	* Must quit job to receive benefits
Contribution requirement	* None	* None
Post reform		
Eligibility status	* All men, regardless of household status * Age 60 and above	* All men, regardless of household status * Age 70 and above, or 30 years of work
Benefit amount	* 100% of minimum wage	* Min. benefit: 100% of minimum wage * Rises with work and earnings
Work requirement	* Min. work requirements increase from 5 years in 1991 to 15 years in 2011	* No requirement to quit job to receive benefits
Contribution requirement	None	* Working wages are taxed

A Model Appendix: Derivations

A.1 Life-cycle model of labor supply (comparing steady states)

The model begins by describing how people choose the number of years they plan to work over their lifetime. Assume that individual i from cohort c lives $\bar{a}_c$ years and receives utility from consumption of market goods, C, and of home goods, H, over their lifetime. Their consumption of home goods is inversely proportional to the number of years the individual spends on market work, L, with $H = \bar{a}_c - L$. We will consider L the individual's choice variable. Without the pension regime, individuals receive a market wage, w, that does not change over their lifetime. As we are modeling decisions of secondary earners, we assume the household has access to non-pension wealth W, with or without the reform.

Under the pension regime, the individual receives a pension with present discounted value of $\tilde{P}$ if they work at least $\tilde{L}_c$ years over their lifetime.[15] Assume that under the pension regime, people receive a market wage w_P, which may or may not be the market wage that prevails without the pension.

The optimization problem is then:

$$\begin{aligned} &\max_{H} \ (1-\alpha_i)log\mathrm{C} \ + \alpha_i log\mathrm{H} \\ s.t. \quad &C = \sum_t c_t, \quad L = \sum_t \ell_t, \quad H = \bar{a}_c - L \\ &C \leq \begin{cases} w(\bar{a}_c - H) + W & \text{without pension} \\ w_p(\bar{a} - H) + W + \tilde{P} \times \mathbb{1}_{L \geq \tilde{L}_c} & \text{with pension} \end{cases} \end{aligned}$$

The value of home goods, which also can be interpreted as the opportunity cost of market work, is captured by α_i and can be heterogeneous across people.

Any pension scheme, $\{\tilde{P}, \tilde{L}_C\}$, affects individuals differently according to their opportunity cost of market work, α_i. The solution, detailed below identifies three types of workers: those who work regardless of whether the pension is offered (market workers), those who work exactly the number of years the pension requires for eligibility (compliers), and those who never worker regardless of whether the pension is offered (non-market workers). Figure 5 describes how these groups vary according to their opportunity cost of market work, with an assumed lifespan of 80 years and under the introduction of a pension scheme that required 15 years of work.

A.1.1 Interior solutions

The interior solution to the individual optimization problem without a pension is:

$$\begin{aligned} C^*_{NP} =& (1-\alpha_i)(w\bar{a}_c + W) \\ L^*_{NP} =& (1-\alpha_i)(\bar{a}_c + W) \end{aligned} \tag{3}$$

[15] Note this abstracts from the eligibility age for now. Since we assume zero discounting, people are indifferent regarding which year in their life they work, so they will work before any eligibility age to maximize the present discounted value of their pension.

The interior solution under the pension regime is:

$$\begin{aligned} L_P^* =& (1-\alpha_i)(\bar{a}_c + W) - \frac{\alpha_i \tilde{P}}{w_P} \\ C_P^* =& (1-\alpha_i)(w_p \bar{a}_c + W + \tilde{P}) \end{aligned} \tag{4}$$

A.1.2 Types of workers

The group identified as "market workers" will choose the interior solution under both the no-pension and the pension regime, L_{NP}^* and L_P^*, respectively. The group defined as "compliers" will choose to work at the interior solution, L_{NP}^*, under the no-pension regime and exactly $\tilde{L}_c \neq L_P^*$ years under the pension regime. The group defined as non-market workers will work L_{NP}^* years regardless of whether a pension regime is in place.

Market workers are people with a low value of home production who choose to work more than $\tilde{L}_c$ when pensions are offered: $L_P^* \geq \tilde{L}_c$, or $(1-\alpha_i)(\bar{a}_c + W) - \frac{\alpha_i \tilde{P}}{w_P} \geq \tilde{L}_c$. This is true for all individuals with preferences such that $\alpha_i \leq \frac{w_p(\bar{a}_c + W - \tilde{L}_c)}{(\bar{a}_c + W)w_P + \tilde{P}} \equiv \alpha_1$.

Compliers are people with a slightly higher value of home production who work exactly the number of years required for pension eligibility. Compliers can be separated into two groups:

- People who work $\tilde{L}_c$ years under the pension regime and *less* than they would have if no pension were offered. This group has preferences such that:

$$\begin{aligned} L_{NP}^* \geq \tilde{L}_c \geq L_P^* \\ (\bar{a}_c + W)(1-\alpha_i) \geq \tilde{L}_c \geq (1-\alpha_i)(\bar{a}_c + W) - \frac{\alpha_i \tilde{P}}{w_P} \\ \bar{a}_c(1-\alpha_i) \geq \tilde{L}_c - W(1-\alpha_i) \geq (1-\alpha_i)\bar{a}_c - \frac{\alpha_i \tilde{P}}{w_P} \end{aligned} \tag{5}$$

  This is true for $\alpha_i \in \left(\frac{w_p(\bar{a}_c + W - \tilde{L}_c)}{(\bar{a}_c + W)w_P + \tilde{P}}, \frac{W + \bar{a}_c - \tilde{L}_c}{W + \bar{a}_c} \right)$. Defined $\alpha_2 = \frac{W + \bar{a}_c - \tilde{L}_c}{W + \bar{a}_c}$.

- People who work $\tilde{L}_c$ years under the pension regime and *more* than they would have if no pension is offered. They do so if the value of working $\tilde{L}_c$ and receiving a pension is larger than the value of working $\tilde{L}_{NP}^*$ and not receiving a pension. This group has preferences such that:

$$L_{NP}^* \leq \tilde{L}_c \tag{6}$$

  and

$$V_P^*(\tilde{L}) > V_{NP}^*(L_{NP}^*) \tag{7}$$

  In the below proposition, we show that for the group of compliers, α_i is bounded above by a finite value α_3, and that $\alpha_3 > \frac{\bar{a}_c - \tilde{L}_c}{\bar{a}_c}$ for any $\tilde{P} > 0$, implying that

compliers value non-market work at level:

$$\alpha_i \in \left(\frac{\bar{a}_c - \tilde{L}}{\bar{a}_c}, \alpha_3 \right] \tag{8}$$

Proposition 1: Existence and Description of Compliers The conditions for an individual to be a complier:

$$L_P^* < \tilde{L}_c \iff \alpha_i > \frac{\bar{a}_c - \tilde{L}}{\bar{a}_c} \quad AND \quad V_P^*(\tilde{L}) > V_{NP}^*(L_{NP}^*) \tag{9}$$

The constraint $V_P^*(\tilde{L}) > V_{NP}^*(L_{NP}^*)$ stipulates that the individual prefers to engage in more market work and receive the pension, rather than maintaining lower levels of market work and not receiving the pension. We can solve this constraint to find the maximum value of α for which the pension scheme encourages the individual to engage in more market work. We will use the combination of the above two conditions to find this. First, let us document:

$$V_{NP}^*(L_{NP}^*) = (1-\alpha) log\{w(1-\alpha)(\bar{a_c} + W)\} + \alpha log \alpha(\bar{a}_c + W) \tag{10}$$

And

$$V_P^*(\bar{L}) = (1-\alpha) log\{w(\bar{a}_c - \bar{L}) + P + W\} + \alpha log \bar{L} \tag{11}$$

So, for $V_P^*(\bar{L}) > V_{NP}^*(L_{NP}^*)$, we must have

$$(1-\alpha) log \frac{\{w(\bar{a}_c - \bar{L}) + P + W\}}{\{(1-\alpha)(w\bar{a}_c + W)\}} > \alpha log \frac{\alpha \bar{a}_c}{\bar{L}} \tag{12}$$

Or

$$log \frac{\{w(\bar{a}_c - \bar{L}) + P + W\}}{w\bar{a}_c + W} > log(1-\alpha) + \frac{\alpha}{(1-\alpha)} log\left(\frac{\alpha \bar{a}_c}{\bar{L}}\right) \tag{13}$$

The right-hand side of the inequality is monotonically increasing in α, so there exists an $\tilde{\alpha}$ such that compliers exist in the range: $\frac{\bar{L}}{\bar{a}_c} < \alpha < \tilde{\alpha}$.

We can further derive that α is bounded above by a finite value $\tilde{\alpha}$, and that $\tilde{\alpha} > \frac{\bar{L}}{H}$ for any $P > 0$.

Never market workers have a high value of home production: even when a pension is available, they choose to work less than $\tilde{L}_c$ and, therefore, do not receive the pension. That is, they meet two conditions:

$$L_{NP}^* < \tilde{L}_c \iff \alpha_i > \frac{\bar{a}_c - \tilde{L}_c}{\tilde{L}_c} \quad AND \quad V_P^*(\tilde{L}) \leq V_{NP}^*(L_{NP}^*) \tag{14}$$

Accordingly, non-market workers place high value in non-market activity, with $\alpha_i > \alpha_3$

Figure 5 summarizes these results. Let $\alpha_1 = \frac{w(\bar{a}_c + W - \tilde{L}_c)}{w(\bar{a}_c + W) + P}$, $\alpha_2 = \frac{\bar{a}_c - \tilde{L}}{\bar{a}_c}$, and α_3 be the cut-off described in proposition 1. From 0 to α_1, we see people who respond only to the

wealth effect – they work less over their lifetime (unconstrained market workers). Then from α_1 to α_2, both the wealth and the eligibility effect are active (constrained market workers); the wealth effect of the pension encourages them to work less, but the eligibility effect prevents them from decreasing their labor supply below $\tilde{L}_c$. Finally, with $\alpha > \alpha_2$, people are working more to receive the pension (compliers).

A.2 Including the Spousal Pension

Suppose that the pension expansion includes an increase in pension wealth to the individual's partner of $\tilde{P}_S$. In this extension, we abstract from other household wealth. The individual's new optimization problem is:

$$\begin{aligned} &\max_{H} \ (1-\alpha_i) log\mathrm{C} \ + \alpha_i log\mathrm{H} \\ s.t. \quad &C = \sum_t c_t, \quad L = \sum_t \ell_t, \quad H = \bar{a}_c - L \\ &C \leq \begin{cases} w(\bar{a}_c - H) & \text{without pension} \\ w_p(\bar{a} - H) + \tilde{P}_S + \tilde{P} \times \mathbb{1}_{L \geq \tilde{L}_c} & \text{with pension} \end{cases} \end{aligned}$$

The interior solution with the pension expansion to an individual who chooses not to work the minimum number of years required to receive their own pension is:

$$\begin{aligned} C^*_{SNP} =& w_p(1-\alpha_i)(\bar{a}_c + \tilde{P}_S) \\ L^*_{SNP} =& (1-\alpha_i)\bar{a}_c - \frac{\alpha_i \tilde{P}_S}{w_p} \end{aligned} \tag{15}$$

The interior solution for the individual who works the required number of years to receive their own pension is:

$$\begin{aligned} L^*_{SP} =& (1-\alpha_i)\bar{a}_c - \frac{\alpha_i(\tilde{P} + \tilde{P}_S)}{w_P} \\ C^*_{SP} =& w_p(1-\alpha_i)(\bar{a}_c + \tilde{P} + \tilde{P}_S) \end{aligned} \tag{16}$$

Types of workers Introducing the spousal pension implies that all workers with affected spouses will work less due to the wealth effect of their partner's pension. The thresholds between market workers, up-compliers, and down-compliers, however, will be affected by the increase in spousal pension wealth.

With an increase in spousal pensions, the only non-responders will be some who previously did not work at all. Instead, some of the previous non-responders with high αs will work less, but not enough to become eligible for their own pension.

Market workers are those who would work more than the required $\tilde{L}$ years with a spousal pension: $L^*_{SP} > \tilde{L}$, which occurs when:

$$\alpha_i < \frac{w_p(\bar{a}_c - \tilde{L})}{w_p\bar{a} + \tilde{P} + \tilde{P}_S} = \alpha_1^S \tag{17}$$

Compliers work less that they would without the expansion: $L^*_{SP} < \tilde{L}$ and $V^*_{SP}(\tilde{L}) > V^*_{SNP}(L^*_{SNP})$. Whether they are down- or up-compliers depends on whether $\tilde{L}$ is higher or lower than L^*_{NP}.

Both market and home workers If people worked at all prior to the pension, they either increased labor supply over their lifetime to comply with the work requirement, or they decreased labor supply due to a wealth effect from their partners pension. If non-responders exist they are not in the workforce at all with or without the pension expansion.

A.3 Allowing for Part-time Work

Intuitively, we would expect the constrained workers (both up and down-compliers who respond to the eligibility effect) to choose part time work if it is available. Let $\kappa < 1$. A worker who chooses part-time work for one year works $\kappa\ell$ in that year. If they chose part-time work throughout their lifetime then lifetime labor supply is κL, and lifetime home-good consumption is $H = \bar{a} - \kappa L$. In this extension, we abstract from other household wealth.

$$\begin{aligned} &\max_H \ (1-\alpha_i) log\mathrm{C} \ + \alpha_i log H \\ s.t. \quad &C = \sum_t c_t, \quad L = \sum_t \ell_t, \quad H = \bar{a}_c - L \\ &C \leq \begin{cases} \kappa w(\bar{a}_c - H) & \text{without pension} \\ \kappa w_p(\bar{a} - H) + \tilde{P} \times \mathbb{1}_{L \geq \tilde{L}_c} & \text{with pension} \end{cases} \end{aligned}$$

Without a pension, κ and L are not separable, so optimal lifetime labor supply is:

$$\begin{aligned} \kappa L &= (1-\alpha_i)\bar{a} \\ C &= (1-\alpha_i) w \bar{a} \end{aligned}$$

With a pension and no work requirement, κ and L are similarly non-separable:

$$\begin{aligned} \kappa L &= (1-\alpha_i)\bar{a} - \frac{\alpha_i P}{w_P} \\ C &= (1-\alpha_i) w_p \bar{a} + (1-\alpha_i) P \end{aligned}$$

With a work requirement, the individual receives the full pension P only if they chose $L > \tilde{L}$. We could represent this choice graphically by re-interpreting the x-axis on our lifetime labor supply graph as $\frac{\alpha_i}{\kappa}$ or adding a z-axis for two dimensional preferences: α_i and κ_i. This would increase the portion of individuals who are "up-compliers:" if they have a particularly high α, they could work part-time.

A.4 "Decomposing" the diff-in-diff estimates by cohort

This section expands the static model of lifetime labor supply to predict people's labor supply responses to a pension expansion that occurs during their working lives. In so doing, we model the average annual treatment effect of pension expansion on extensive-margin labor supply within a cohort and explore the channels through which this effect operates. In this step, we add the assumption that people are eligible to receive the

rural pension, provided that they have worked at least $\tilde{L}_c$ years, only when they reach eligibility age, $\tilde{a}_E$. The pension scheme is now described by the triple $\mathcal{P} = \{\tilde{L}, \tilde{P}, \tilde{a}_E\}$.

We assume that each member of cohort c has the same target retirement age and that $\alpha_i \sim G(\alpha)$ within a cohort, implying that the aggregate lifetime labor supply of cohort c under the no-pension regime is $L_C^{NP} = \int_\alpha L_{NP}^* dG(\alpha)$, where L_{NP}^* is the interior solution for lifetime labor supply under the no-pension regime in Section A.1. At the cohort level, aggregate lifetime labor supply is smoothed across the years before the cohort reaches its target retirement date. Accordingly, the extensive-margin labor supply within a cohort in a given year is $\int_\alpha \frac{L_C^{NP}}{\bar{a}_c} dG(\alpha)$ if everyone within a cohort planned to retire only upon death (ie., the cohort's retirement age is $\bar{a}_c$).

Let $\bar{a}_R$ be the cohort's target retirement age under the no-pension regime, and $\tilde{a}_R$ be the cohort's target retirement age under the pension regime. We allow the retirement age to differ under the pension regime, despite the fact that there is no explicit incentive from a lifetime perspective for people to adjust their personal ideal retirement age simply because the statutory pension eligibility age has decreased.[16] Assuming that people may begin work as soon as they are "born," the possible working life is $\bar{a}_R$ years under the no-pension regime and $\tilde{a}_R$ years under the pension regime.

A pension scheme $\mathcal{P}$ is introduced in year j, when cohort c is age $a_{cj} = j - c$ years of age. Prior to the pension introduction, cohort c planned to work $\int_\alpha L_{NP}^* dG(\alpha)$ years before retirement. At the cohort level, these years are smoothed out over the working life so that the cohort works $\frac{L_{NP}^*}{\bar{a}_R}$ per year. By year j, cohort c has worked $\frac{a_{cj}}{\bar{a}_R} \int_\alpha L_{NP}^* dG(\alpha)$ years. Had the pension not been introduced, cohort c would have continued to work $\frac{1}{\bar{a}_R} \int_\alpha L_{NP}^* dG(\alpha)$ per year until age $\bar{a}_R$. Define L_{ct}^{NP} to be cohort c's labor supply in period $t > j$ if the pension were never introduced:

$$L_{ct}^{NP} = \frac{1}{\bar{a}_R} \int_\alpha L_{NP}^* dG(\alpha) \tag{18}$$

We define the average annual treatment effect of the pension on cohort c to be $\Delta L_{ct} = L_{ct}^P - L_{ct}^{NP}$, where L_{ct}^{NP} is as defined in equation 18 and L_{ct}^P is the labor supply of cohort c in year $t > j$ after the pension has been introduced. The next step is to find an expression for L_{ct}^P. Under the pension regime, suppose that individual i of cohort c plans to work $L_P(\alpha_i)$ years before age $\tilde{a}_R$. As of time j, they have already worked $\frac{a_{cj}}{\bar{a}_R} L_{NP}^*$ years. Accordingly, they must now work an additional $L_P(\alpha_i) - \frac{a_{cj}}{\bar{a}_R} L_{NP}^*$ years before age $\tilde{a}_R$, where $L_P(\alpha_i)$ is their desired lifetime labor supply under the pension regime. At the cohort level, these years of work will be smoothed over the rest of the cohort's working life, which is $\tilde{a}_R - a_{cj}$ additional years. Therefore, the labor supply of cohort c in year $t > j$ after the pension has been introduced is:

$$L_{ct}^P = \frac{1}{\tilde{a}_R - a_{cj}} \int_\alpha L_P(\alpha_i) - \frac{a_{cj}}{\bar{a}_R} L_{NP}^* dG(\alpha)$$

The average treatment effect on extensive-margin labor supply in each year $t > j$ is $\Delta L_{ct} = L_{ct}^P - L_{ct}^{NP}$, or:

[16] Several papers, including Behaghel & Blau (2012) and Neumark & Song (2013), show discontinuously higher retirement rates at pension elibigibility. Some explanations include credit constraints. We show in appendix B.2 that a discontinuity develops for women at age 55 and for men at age 60 following this change.

$$
\begin{aligned}
\Delta L_{ct} =& L_{ct}^{P} - L_{ct}^{NP} \\
\Delta L_{ct} =& \underbrace{\frac{1}{\tilde{a}_R - a_{cj}} \int_{\alpha} L_P(\alpha_i) - \frac{a_{cj}}{\bar{a}_R} L_{NP}^{*} dG(\alpha)}_{L_{ct}^{P}} - \underbrace{\frac{1}{\bar{a}_R} \int_{\alpha} L_{NP}^{*} dG(\alpha)}_{L_{NP}^{*}} \\
=& \int_{\alpha} \frac{1}{\tilde{a}_R - a_{cj}} L_P(\alpha_i) - L_{NP}^{*} \left(\frac{a_{cj}}{\bar{a}_R} \frac{1}{\tilde{a}_R - a_{cj}} + \frac{1}{\bar{a}_R} \right) dG(\alpha) \\
=& \int_{\alpha} \frac{1}{\tilde{a}_R - a_{cj}} L_P(\alpha_i) - L_{NP}^{*} \left(\frac{a_{cj} + \tilde{a}_R - a_{cj}}{\bar{a}_R(\tilde{a}_R - a_{cj})} \right) dG(\alpha)
\end{aligned}
$$

$$
\Delta L_{ct} = \int_{\alpha} \frac{1}{\tilde{a}_R - a_{cj}} L_P(\alpha_i) - L_{NP}^{*} \left(\frac{\tilde{a}_R}{\bar{a}_R(\tilde{a}_R - a_{cj})} \right) dG(\alpha) \tag{19}
$$

A.4.1 Decomposing into wealth and retirement timing effect with no work requirement

Consider, first, a pension scheme without a work requirement, $\tilde{L}_c = 0$. In this case, $L_P(\alpha_i) = L_P^{*} = (\bar{a}_c + W)(1 - \alpha_i) - \frac{\alpha_i \tilde{P}}{w_P}$ and $L_{NP}^{*} = (1 - \alpha_i)(\bar{a}_c + W)$.The difference-in-differences estimate from equation 19 is:

$$
\Delta L_{ct} = \int_{\alpha} \frac{1}{\tilde{a}_R - a_{cj}} \underbrace{(\bar{a}_c + W)(1 - \alpha_i) - \frac{\alpha_i \tilde{P}}{w_P}}_{L_P = L_P^{*}} - \underbrace{(1 - \alpha_i)(\bar{a}_c + W)}_{L_{NP}^{*}} \left(\frac{\tilde{a}_R}{\bar{a}_R(\tilde{a}_R - a_{cj})} \right) dG(\alpha)
$$

$$
\Delta L_{ct} = \int_{\alpha} (1 - \alpha_i)(\bar{a}_c + W) \left(\frac{1}{\tilde{a}_R - a_{cj}} - \frac{\tilde{a}_R}{\bar{a}_R(\tilde{a}_R - a_{cj})} \right) - \frac{1}{\tilde{a}_R - a_{cj}} \frac{\alpha_i \tilde{P}}{w_P} dG(\alpha)
$$

$$
\Delta L_{ct} = \int_{\alpha} (1 - \alpha_i)(\bar{a}_c + W) \left(\frac{\bar{a}_R - \tilde{a_R}}{\bar{a}_R(\tilde{a}_R - a_{cj})} \right) - \frac{1}{\tilde{a}_R - a_{cj}} \frac{\alpha_i \tilde{P}}{w_P} dG(\alpha)
$$

The average annual treatment effect by cohort from equation 19 can then be decomposed into a retirement-timing effect and a wealth effect:

$$
\Delta L_{ct} = \frac{1}{\tilde{a}_R - a_{cj}} \int_{\alpha} \underbrace{(1 - \alpha_i)(\bar{a}_c + W) \left(\frac{\bar{a}_R - \tilde{a_R}}{\bar{a}_R} \right)}_{\text{Retirement Timing Effect}} \underbrace{- \frac{\alpha_i \tilde{P}}{w_P}}_{\text{Wealth Effect}} dG(\alpha) \tag{20}
$$

A pension scheme without a work requirement influences the average annual treatment effect of the extensive-margin labor supply response through two channels. First, labor supply decreases through the wealth effect by an amount that depends on the productivity of home-work, the lifetime value of the pension, and the wage the individual can receive under the pension regime. Second, years of labor may be redistributed over the lifetime through the retirement timing effect. A pension scheme that encourages people to decrease their target retirement age, $\bar{a}_R > \tilde{a}_R$, imposes a positive retirement timing effect until age $\tilde{a}_R$ and negative effect from age $\tilde{a}_R$ to $\bar{a}_R$. A pension scheme that encourages people to increase their target retirement age, $\tilde{a}_R > \bar{a}_R$, imposes a negative retirement-timing effect until age $\bar{a}_R$ and positive from age $\bar{a}_R$ to $\tilde{a}_R$. Both the wealth

effect and retirement-timing effect are larger in magnitude if the cohort is closer to retirement age when the pension reform occurs; the average annual treatment effect by cohort is closer to zero for cohorts that are younger when the reform was enacted.

A.4.2 Including the work requirement

Next, consider a pension expansion with a work requirement: people in cohort c must work for at least $\tilde{L}_c$ years to receive pension eligibility. Introducing this requirement creates three additional groups of workers. Two of these groups comply with the work requirement by working exactly $\tilde{L}_c$ years in their life: up-compliers work more than they would have without the pension, and down-compliers work less than they would have without the pension. The third group of workers does not find the pension benefit large enough to motivate them to work the required number of years; these non-responders choose to forgo the pension and, instead, work the same number of years they would have if the pension had not been introduced.

Down- and Up-Compliers work $\tilde{L}_c$ years under the pension scheme. We can write the difference-in-difference estimate for this population with $\alpha_i \in (\alpha_1, \alpha_3]$, by setting $L_P^* = \tilde{L}$ in equation 19. Note that $\tilde{L}_c = \tilde{L}_c + L_P^* - L_P^*$, and equation 19 for the group of compliers becomes:

$$\Delta L_{ct} = \int_{\alpha_1}^{\alpha_2} \frac{1}{\tilde{a}_R - a_{cj}} (\tilde{L}_c + L_P^* - L_P^*) - L_{NP}^* \left(\frac{\tilde{a}_R}{\bar{a}_R(\tilde{a}_R - a_{cj})} \right) dG(\alpha)$$

$$\Delta L_{ct} = \frac{1}{\tilde{a}_R - a_{cj}} \int_{\alpha_1}^{\alpha_2} \underbrace{L_P^* - L_{NP}^* \frac{\tilde{a}_R}{\bar{a}_R}}_{\mathcal{A}} + (\tilde{L}_c - L_P^*) dG(\alpha)$$

Note that we calculated $\mathcal{A}$ in our derivation of the change in labor supply without a work requirement. Accordingly, we see that the retirement-timing effect and the wealth effect are both active for compliers as well as market workers:

$$\Delta L_{ct} = \frac{1}{\tilde{a}_R - a_{cj}} \int_{\alpha_1}^{\alpha_2} \underbrace{\underbrace{(1-\alpha_i)(\bar{a}_c + W)\left(\frac{\bar{a}_R - \tilde{a_R}}{\bar{a}_R}\right)}_{\text{Retirement Timing Effect}} - \underbrace{\frac{\alpha_i \tilde{P}}{w_P}}_{\text{Wealth Effect}}}_{\mathcal{A}} + \underbrace{(\tilde{L}_c - L_P^*)}_{\text{Eligibility Effect}} dG(\alpha)$$

The difference-in-difference estimate of compliers is influenced by a third channel, which we call the eligibility effect. Without the work requirement, compliers would have worked less than $\tilde{L}_c$ when the pension was introduced. However, these workers find the value of receiving the pension to be high enough that they are willing to work the required number of years to achieve eligibility. Thus they work an additional $\tilde{L}_c - L_P^*$ years more than the "interior solution" under an analogous pension regime with no work requirement. Plugging in the equation for L_P^*, we have the average annual treatment effect by cohort for compliers with $\alpha_i \in (\alpha_1, \alpha_3]$:

$$\Delta L_{ct} = \frac{1}{\tilde{a}_R - a_{cj}} \int_{\alpha_1}^{\alpha_2} \underbrace{(1-\alpha_i)(\bar{a}_c + W)\left(\frac{\bar{a}_R - \tilde{a_R}}{\bar{a}_R}\right)}_{\text{Retirement Timing Effect}} - \underbrace{\frac{\alpha_i \tilde{P}}{w_P}}_{\text{Wealth Effect}} + \underbrace{\tilde{L}_c - (1-\alpha_i)(\bar{a}_c + W) + \frac{\alpha_i \tilde{P}}{w_P}}_{\text{Eligibility Effect}} dG(\alpha) \tag{21}$$

Non-responders do not adjust their lifetime labor supply when the pension regime is introduced. Non-responders who do not work at all do not contribute to the cohort-level average annual treatment effect. However, non-responders who work some small amount over their lifetime and adjust their retirement age along with the rest of their cohort would contribute to the retirement-timing effect:

$$\Delta L_{ct} = \frac{1}{\tilde{a}_R - a_{cj}} \int_{\alpha_3}^{1} (1-\alpha_i)(\bar{a}_c + W)\left(\frac{\bar{a}_R - \tilde{a}_R}{\bar{a}_R}\right) dG(\alpha) \tag{22}$$

A.4.3 Complete average annual treatment effect

Accordingly, all members of the cohort respond with the retirement timing effect, everyone but non-responders are influenced by the wealth effect, and only the compliers respond to the eligibility effect. The average annual treatment effect for the whole cohort, including all three groups of workers, is:

$$\Delta L_{ct} = \frac{1}{\tilde{a}_R - a_{cj}} \left[\int_{\alpha} \underbrace{(1-\alpha_i)(\bar{a}_c + W)\left(\frac{\bar{a}_R - \tilde{a_R}}{\bar{a}_R}\right)}_{\text{Retirement Timing Effect}} - \underbrace{\frac{\alpha_i \tilde{P}}{w_P}}_{\text{Wealth Effect}} \; dG(\alpha) + \int_{\alpha_1}^{\alpha_3} \underbrace{\tilde{L}_c - (1-\alpha_i)(\bar{a}_c + W) + \frac{\alpha_i \tilde{P}}{w_P}}_{\text{Eligibility Effect}} dG(\alpha) \right] \tag{23}$$

B Empirical Results, Robustness

B.1 Robustness: Difference-in-differences

This section includes several robustness checks of our difference-in-differences estimates. First, it includes an indicator of whether the spouse is present in the household and the number of household members as the controls in the main difference-in-differences specification. Second, it controls for potential differences in age composition across rural and urban areas by including age and age squared in the control vector. Third, it includes controls for state-by-year and state-by-rural fixed effects. Fourth, it runs the main difference-in-differences specifications comparing rural married women to two alternate control groups: rural men and rural single women.

Figure A.1: Pension and Work Status in Rural versus Urban Areas among Married Women, Difference-in-Difference Estimates, Controlled regressions

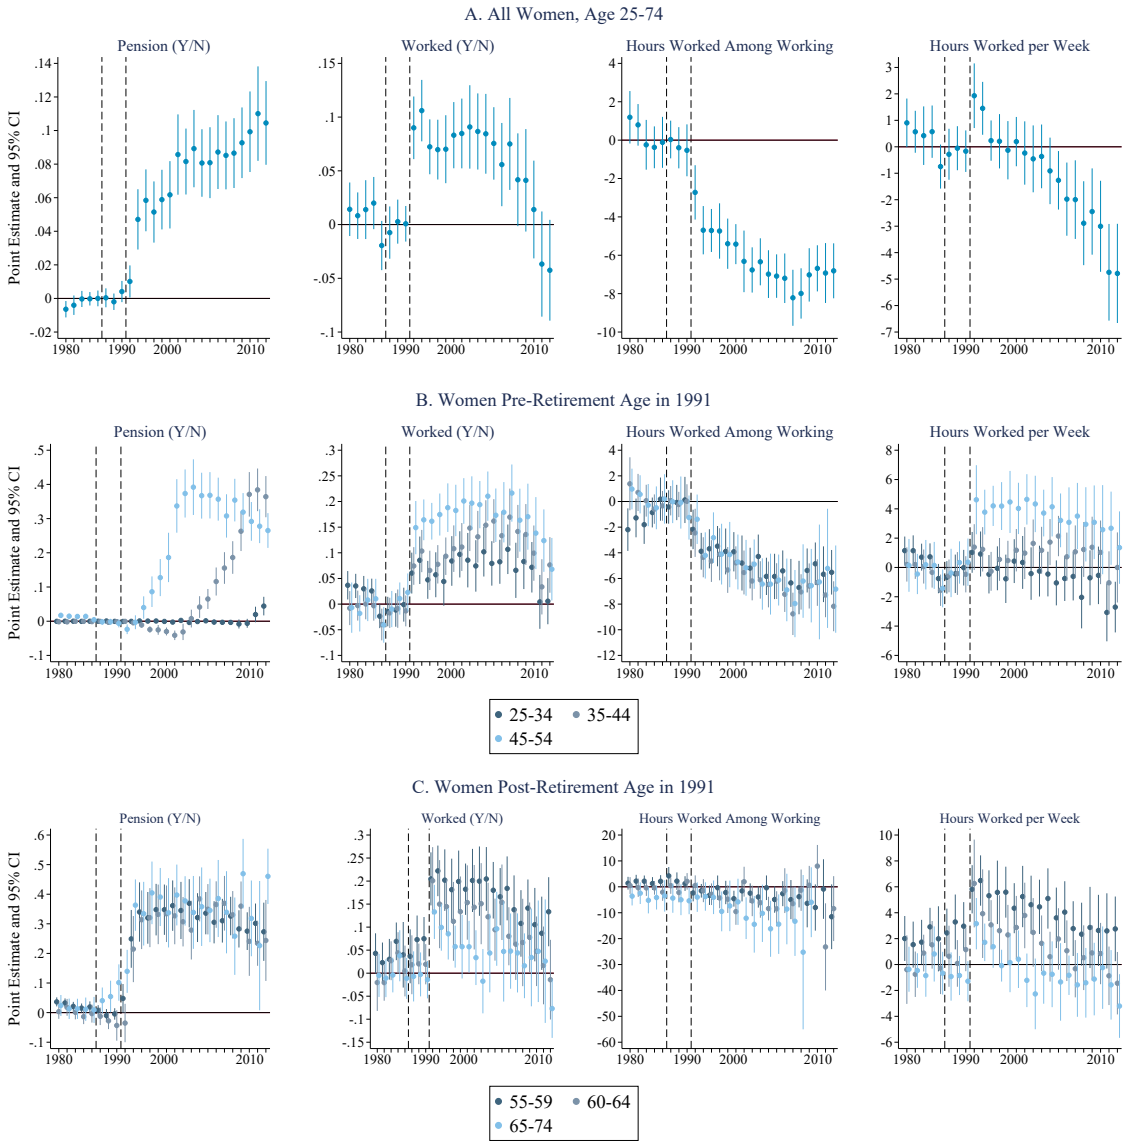

Notes. This Figure is analogous to Figure 2 in the main text, but controls for whether the spouse is present in the household and the number of other household members, $\Gamma'_{ist}X_{ist}$ in the difference-in-difference specification. Each panel shows the β coefficient estimates and 95% confidence intervals on each year from an extended difference in difference regression of the form $y_{ist} = \alpha \times \text{Rural}_{ist} + \sum_{j=1981}^{1988} \beta_j^{pre} \times \text{Rural}_{isj} + \sum_{j=1989}^{2013} \beta_j^{post} \times \text{Rural}_{isj} + \delta_t + \mu_s + \Gamma'_{ist}X_{ist}$, where y_{it} is the outcome variable of interest and Rural = 1 if the individual lives in a rural area. Unless otherwise noted, panel A includes all married women age 25-74 within the year plotted. Panel B includes three different cohorts of married women who were younger than the retirement age of 55 when the law was passed in 1991. Panel C includes three different cohorts of married women who were older than the retirement age of 55 when the law passed in 1991. Coefficients are estimated relative to 1987, the year before the constitutional amendment announcing expansion of the rural pension scheme. The title of each graph refers to the outcome variable used to generate the difference-in-differences estimates. "Receiving Pension (Y/N)" describes an indicator variable equal to one if the individual is receiving pension payments. "Worked (Y/N)" describes an indicator variable equal to one if the individual reports working in the reference week. "Hours Worked among Working" sets the outcome variable as hours worked per week, but limits the sample to individuals who worked in the reference week. "Hours worked per week" sets the outcome variable as hours worked per week.

Figure A.2: Pension and Work Status in Rural versus Urban Areas among Married Women, Difference-in-Difference Estimates

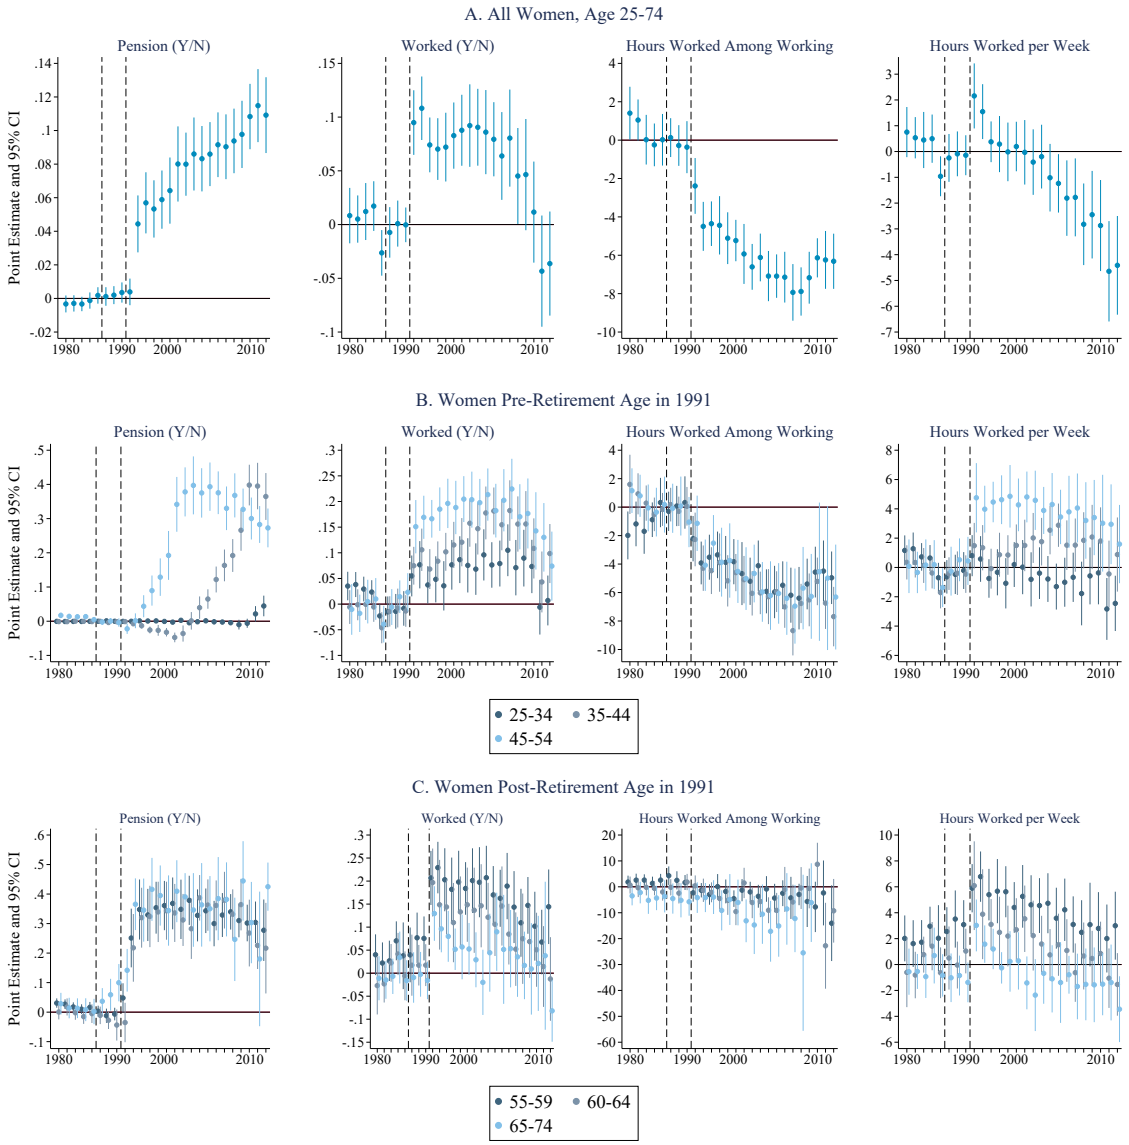

Notes. This Figure is analogous to Figure 2 in the main text, but controls for potential differences in age composition across urban and rural areas with age and age squared, $\Gamma'_{ist}X_{ist}$ in the difference-in-difference specification. Each panel shows the β coefficient estimates and 95% confidence intervals on each year from an extended difference in difference regression of the form $y_{ist} = \alpha \times \text{Rural}_{ist} + \sum_{j=1981}^{1988} \beta_j^{pre} \times \text{Rural}_{isj} + \sum_{j=1989}^{2013} \beta_j^{post} \times \text{Rural}_{isj} + \delta_t + \mu_s + \Gamma'_{ist}X_{ist}$, where y_{it} is the outcome variable of interest and Rural = 1 if the individual lives in a rural area. Unless otherwise noted, panel A includes all married women age 25-74 within the year plotted. Panel B includes three different cohorts of married women who were younger than the retirement age of 55 when the law was passed in 1991. Panel C includes three different cohorts of married women who were older than the retirement age of 55 when the law passed in 1991. Coefficients are estimated relative to 1987, the year before the constitutional amendment announcing expansion of the rural pension scheme. The title of each graph refers to the outcome variable used to generate the difference-in-differences estimates. "Receiving Pension (Y/N)" describes an indicator variable equal to one if the individual is receiving pension payments. "Worked (Y/N)" describes an indicator variable equal to one if the individual reports working in the reference week. "Hours Worked among Working" sets the outcome variable as hours worked per week, but limits the sample to individuals who worked in the reference week. "Hours worked per week" sets the outcome variable as hours worked per week.

Figure A.3: State-by-year, state-by-rural fixed effect: Pension and Work Status in Rural versus Urban Areas among Married Women, Difference-in-Difference Estimates, Controlled regressions

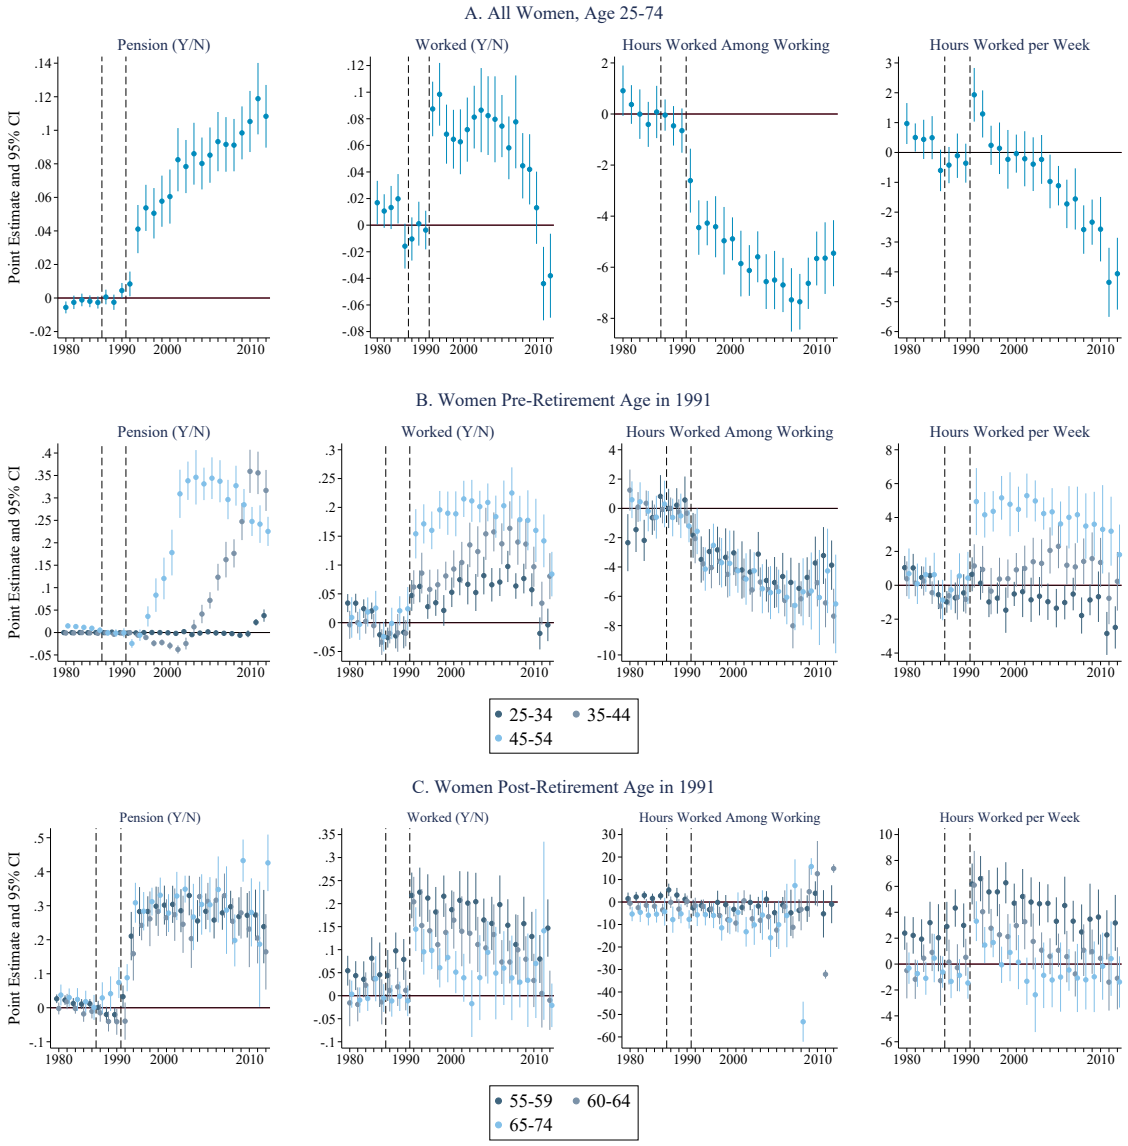

Notes. This Figure is analogous to Figure 2 in the main text, but controls for potential differences in age composition acros urban and rural areas with age and age squared, $\Gamma'_{ist}X_{ist}$ in the difference-in-difference specification. Each panel shows the β coefficient estimates and 95% confidence intervals on each year from an extended difference in difference regression of the form $y_{ist} = \alpha \times \text{Rural}_{ist} + \sum_{j=1981}^{1988} \beta_j^{pre} \times \text{Rural}_{isj} + \sum_{j=1989}^{2013} \beta_j^{post} \times \text{Rural}_{isj} + \delta_t + \mu_s + \Gamma'_{ist}X_{ist}$, where y_{it} is the outcome variable of interest and Rural = 1 if the individual lives in a rural area. Unless otherwise noted, panel A includes all married women age 25-74 within the year plotted. Panel B includes three different cohorts of married women who were younger than the retirement age of 55 when the law was passed in 1991. Panel C includes three different cohorts of married women who were older than the retirement age of 55 when the law passed in 1991. Coefficients are estimated relative to 1987, the year before the constitutional amendment announcing expansion of the rural pension scheme. The title of each graph refers to the outcome variable used to generate the difference-in-differences estimates. "Receiving Pension (Y/N)" describes an indicator variable equal to one if the individual is receiving pension payments. "Worked (Y/N)" describes an indicator variable equal to one if the individual reports working in the reference week. "Hours Worked among Working" sets the outcome variable as hours worked per week, but limits the sample to individuals who worked in the reference week. "Hours worked per week" sets the outcome variable as hours worked per week.

Figure A.4: Difference-in-Difference Estimates, Alternative control groups

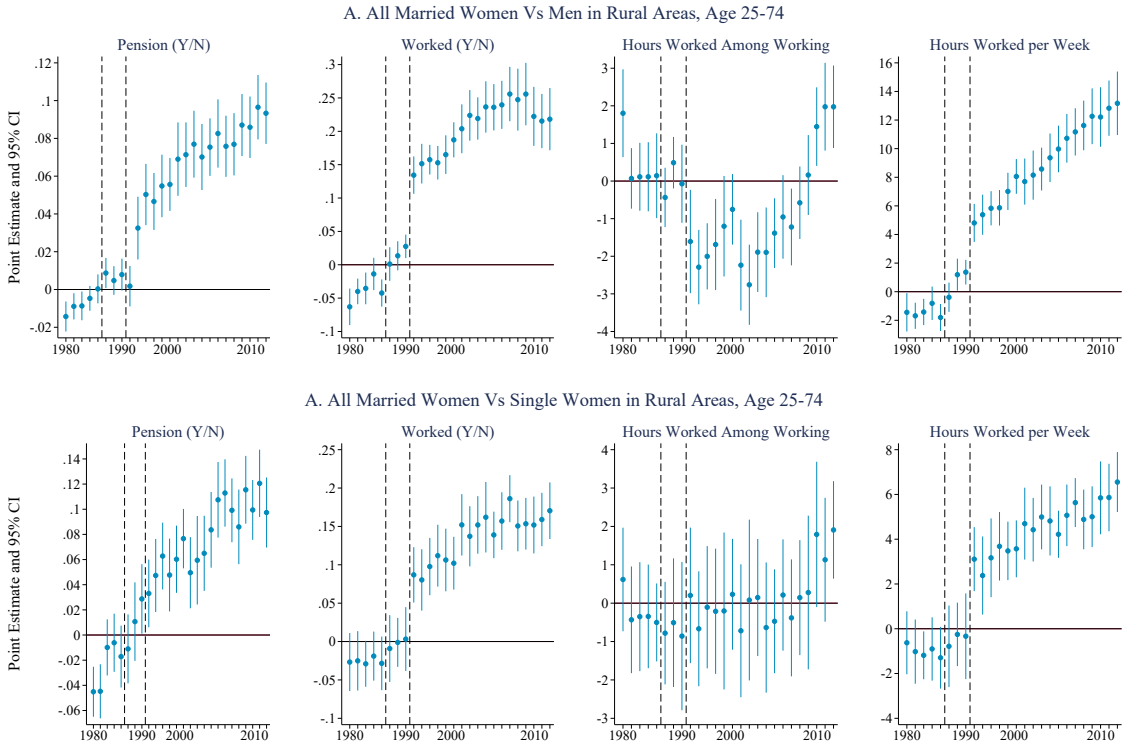

Notes. This figure is analogous to Figure 2 in that it shows difference-in-differences estimates that estimate the impact of the pension expansion on rural married women's labor supply, except it uses alternative control groups. Instead of comparing married rural women to married urban women, Panel A compares married rural women to rural men, and Panel B compares married rural women to single rural women. Prior to the rural pension expansion, rural men and rural single women had access to the rural pension, unlike married rural women. The rural pension expansion doubled the size of the pension for rural men and rural single women; decreased their eligibility age to 60 and 55, respectively; and established the 15-year work requirement that was phased in gradually.

B.2 Difference-in-discontinuities: Robustness

In Appendix B.2, we test the robustness of this discontinuity by running analogous specifications using ages other than the age of eligibility – 50, 60, and 65 – and find no evidence of a discontinuous decrease in labor supply at those ages in any year between 1981 and 2006.

Figure A.5: Difference in Discontinuity at Age 50 (Women)

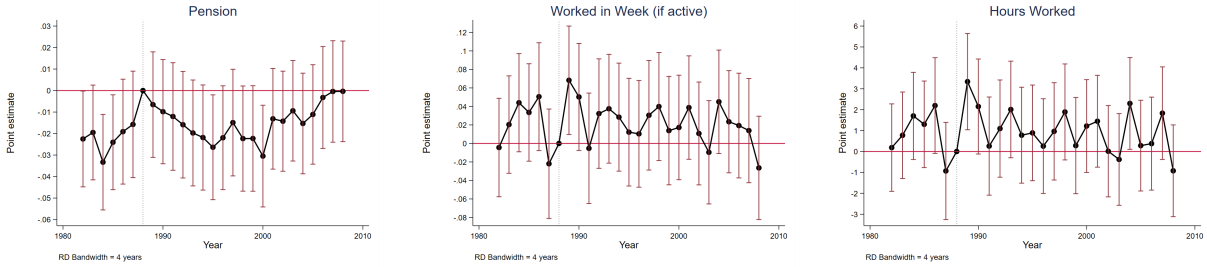

Notes. These graphs show the difference in discontinuity at age 50 in the three variables listed, using a bandwidth of 4 years.Figure shows the discontinuity estimate at age 50 and 95% confidence intervals for RDs run in each year, using a bandwidth of 4 years. Sample is restricted to married women living in rural areas.

Figure A.6: Difference in Discontinuity at Age 60 (Women)

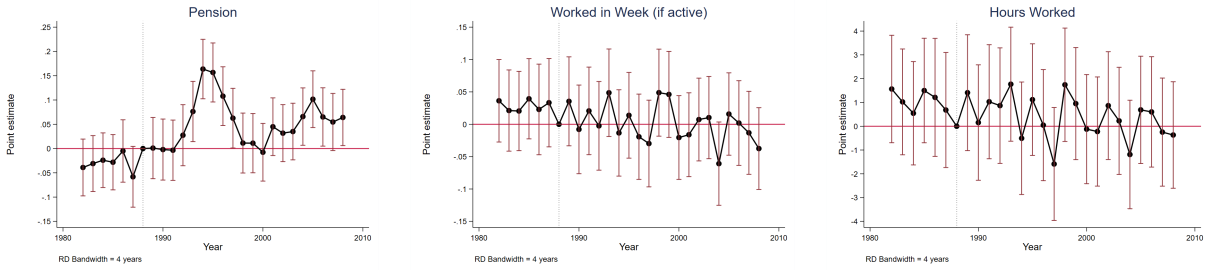

Notes. These graphs show the difference in discontinuity at age 60 in the three variables listed, using a bandwidth of 4 years.Figure shows the discontinuity estimate at age 60 and 95% confidence intervals for RDs run in each year, using a bandwidth of 4 years. Sample is restricted to married women living in rural areas.

Figure A.7: Difference in Discontinuity at Age 65 (Women)

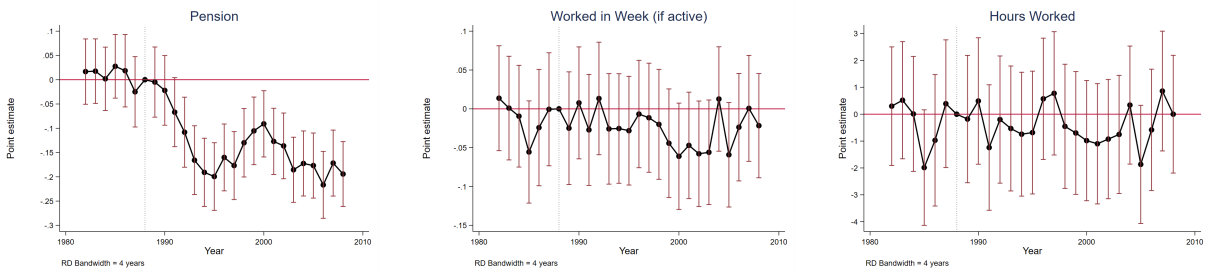

Notes. These graphs show the difference in discontinuity at age 65 in the three variables listed, using a bandwidth of 4 years.Figure shows the discontinuity estimate at age 65 and 95% confidence intervals for RDs run in each year, using a bandwidth of 4 years. Sample is restricted to married women living in rural areas.

Figure A.8: Difference in Discontinuity at Age 55 (Men)

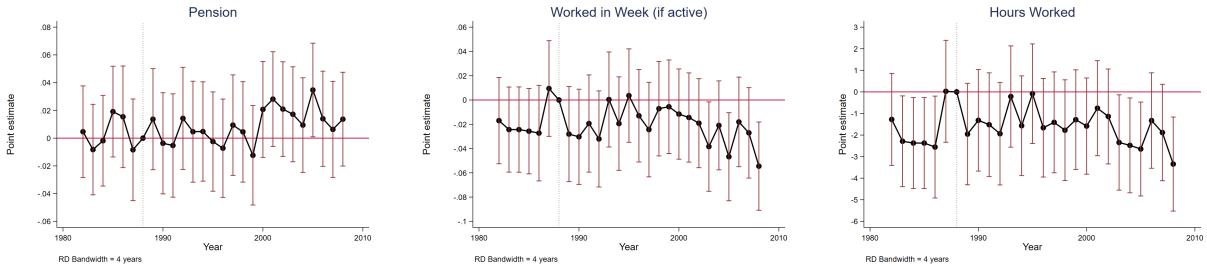

Notes. These graphs show the difference in discontinuity at age 55 in the three variables listed, using a bandwidth of 4 years.Figure shows the discontinuity estimate at age 55 and 95% confidence intervals for RDs run in each year, using a bandwidth of 4 years. Sample is restricted to men living in rural areas.

Figure A.9: Difference in Discontinuity at Age 60 (Men)

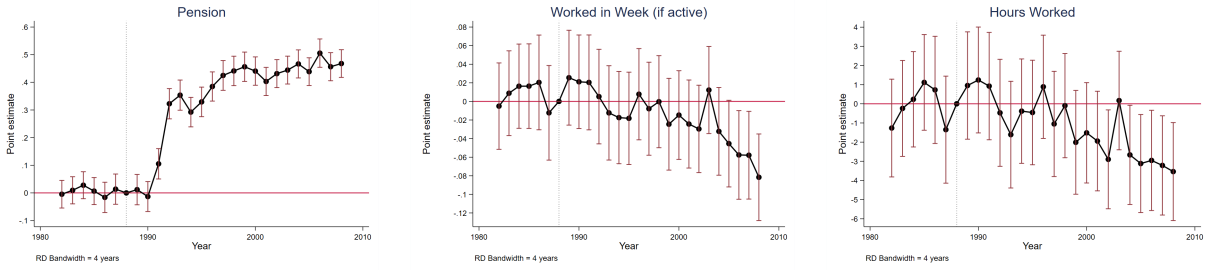

Notes. These graphs show the difference in discontinuity at age 60 in the three variables listed, using a bandwidth of 4 years.Figure shows the discontinuity estimate at age 60 and 95% confidence intervals for RDs run in each year, using a bandwidth of 4 years. The 1991 pension expansion decreased men's eligibility age from 65 to 60. Sample is restricted to men living in rural areas.

Figure A.10: Difference in Discontinuity at Age 65 (Men)

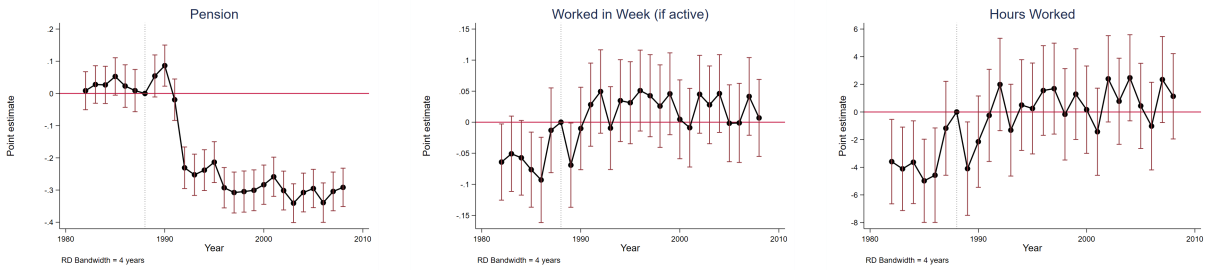

Notes. These graphs show the difference in discontinuity at age 65 in the three variables listed, using a bandwidth of 4 years.Figure shows the discontinuity estimate at age 65 and 95% confidence intervals for RDs run in each year, using a bandwidth of 4 years. The 1991 pension expansion decreased men's eligibility age from 65 to 60. Sample is restricted to men living in rural areas.

C Projecting and Bounding Lifetime Labor Supply Response

Our bounding exercise builds on the following theoretical equation for our annual difference-in-difference estimates:

$$\Delta L_{ct} = \frac{1}{\tilde{a}_R - a_{cj}} \left[\int_{\alpha} \underbrace{(1-\alpha_i)\bar{a}_c \left(\frac{\bar{a}_R - \tilde{a_R}}{\bar{a}_R} \right)}_{\text{Retirement Timing Effect}} \underbrace{- \frac{\alpha_i \tilde{P}}{w_P}}_{\text{Wealth Effect}} \; dG(\alpha) \right. \\ \left. + \int_{\alpha_1}^{\alpha_3} \underbrace{\tilde{L}_c - (1-\alpha_i)\bar{a}_c + \frac{\alpha_i \tilde{P}}{w_P}}_{\text{Eligibility Effect}} dG(\alpha) \right] \tag{24}$$

C.1 Cohorts 45 and younger in 1991

We make two assumptions. First, the pension expansion encouraged these cohorts to decrease their target retirement age: $\tilde{a}_R < \bar{a}_R$. Second, these cohorts decreased their target retirement age to the pension eligibility age: $\tilde{a}_R = \bar{a}_E = 55$. Under these assumptions, the annual change in labor market participation is positive in each year in which the cohort is younger than their adjusted target retirement age, $a_{cj} < \tilde{a}_R$. Moreover, the annual change in labor market participation has the same magnitude for all years in which the cohort is younger than their adjusted target retirement age. This appendix refers to this estimate as $\hat{\beta}$.

Our empirical estimates allow us to evaluate these assumptions and some of their theoretical implications. First, the difference-in-discontinuities estimates, presented in Section 4.3, suggest that a discontinuity in labor force participation developed among cohorts turning 55 in 1998 and later, thus supporting the assumption that $\tilde{a}_R = \bar{a}_E = 55$. Second, Table 2 shows steady increases in labor market participation, between seven and nine percentage points, for the cohort age 25-34 in 1991. This is the only cohort for whom all members had not yet reached the pension eligibility age of 55 by the end of our data. This finding is consistent with the theoretical implication that the annual change in labor market participation has a similar magnitude for all years in which the cohort is younger than their adjusted target retirement age. We assume $\hat{\beta} = .09$ for people between 25 and 34.

Our upper bound estimate of the percentage point increase in labor-force participation for individuals born in year b_c, who are between 25 and 34 in 1991, is then:

$$\hat{\beta} \times (\tilde{a}_R - a_{c,1991}) = .09 \times (55 - (1991 - b_c)) \tag{25}$$

To find the average upper bound for the cohort aged between 25 and 34 in 1991, we then average this upper bound for all individuals born between 1968 and 1957.

For women who were between 35 and 44 in 1991, we again look to Table 2 for the maximum annual increase, and assume that $\hat{\beta} = .16$ for this group. Our upper bound estimate of the percentage point increase in labor-force participation among individuals born in year b_c, who are between 35 and 44 in 1991, is then:

$$\hat{\beta} \times (\tilde{a}_R - a_{c,1991}) = .16 \times (55 - (1991 - b_c)) \tag{26}$$

To find the average upper bound for the cohort aged between 35 and 44 in 1991, we then average this upper bound for all individuals born between 1958 and 1947.

C.2 Cohorts older than 45 in 1991

The target retirement age of older cohorts was potentially limited by the graduated design of the work requirement. Consider a married women, who did not work prior to the pension expansion and was 55 in 1991, that was encouraged by the pension expansion to work the 5 years required to achieve pension eligibility. This 'new market entrant', as Figure 5 identifies her, would have to adjust her target retirement age to 60 or older to provide enough time to accumulate the required years of work. For these older cohorts, we assume a target retirement age equal to the youngest possible age an individual in the cohort could retire and receive the lifetime pension, had they not worked at all prior to the pension expansion. The table below describes our assumed target retirement ages after the expansion, and how they were calculated based on the individual's required years of work to achieve pension eligibility.

Our estimate of $\hat{\beta}$ for these individuals is the highest five-year average estimate in Table 2 for the relevant cohort: $\hat{\beta} = .19$ for people between ages 45 and 55, .17 for people between 55 and 59, and .16 for people between 60 and 69 in 1991. Our upper bound for individuals born in year b_c, who are older than 45 in 1991, is then:

$$\hat{\beta} \times (\tilde{a}_R - a_{c,1991}) \tag{27}$$

where $\hat{\beta}$ and $\tilde{a}_R$ vary by year born.

C.3 Estimates and upper bounds on lifetime labor supply adjustment

Tables A.2 - A.5 present estimates of lifetime adjustments in labor supply for married women and men of various age groups. "Realized adjustments" are those that are aggregated across the years 1992-2012, for which our data is available. "Potential adjustments" calculate lifetime adjustments in labor supply under the two assumptions on $\tilde{a}_R$ described above. The first assumes that the target retirement age adjusts downward to 70 following the reform. The second assumes that the target retirement age adjusts downward to 55, or the earliest possible age at which an individual could begin to work and receive the pension had they not worked prior to the reform.

Table A.2: Lifetime Adjustments in Labor-market Participation: Rural Women

	Age in 1991				
	25-34	35-44	45-54	55-59	60-69
Realized adjustments					
Realized years per worker	3.12	4.07	4.89	3.09	1.94
Pct Increase in worker-years	6.52	8.14	10.49	7.64	4.80
Years per person	1.44	2.15	2.53	1.40	0.72
Potential adjustments (years per person)					
Upper bound, $\tilde{a}_R = 70$	3.64	4.81	3.73	2.02	0.87
Upper bound, $\tilde{a}_R$ set at discontinuity	2.30	2.48	1.43	0.85	0.80

Notes. This table presents estimates of realized and potential lifetime labor supply adjustments for each cohort of rural women, using annual cohort-level coefficients estimated in Table 2. 'Realized' lifetime adjustments refer to increases in labor market participation that occurred and are observed in the data prior to 2012. This includes three measures: 'realized years per worker' aggregates the annual treatment effect for each cohort, as estimated in Table 2, and divides it by the number of workers in that cohort in 1991; 'Percent increase in worker-years' takes the worker-year as the level of analysis and estimates the lifetime treatment effect relative to all years worked by that cohort prior to 2012; and 'Realized years per person' distributes the lifetime treatment effect across the full cohort population in 1991. 'Potential' lifetime adjustment measures project annual adjustments past the available data as described in the text of this appendix. These estimates are comparable to the realized estimates calculated as 'years per person.'

Table A.3: Lifetime Adjustments in Labor-market Participation: Young rural Women

	Age in 1991				
	25-26	27-28	29-30	31-32	33-34
Realized adjustments					
Years per worker	1.62	2.47	3.74	2.85	3.14
Pct Increase in worker-years	4.57	6.34	8.10	5.96	5.74
Years per person	0.88	1.33	1.99	1.43	1.61
Potential adjustments (years per person)					
Upper bound, $\tilde{a}_R = 70$	2.31	3.32	4.29	4.04	3.98
Upper bound, $\tilde{a}_R$ set at discontinuity					

Notes. This table presents estimates of realized and potential lifetime labor supply adjustments for two-year cohorts of rural women who were younger than 35 in 1991. 'Realized' lifetime adjustments refer to increases in labor market participation that occurred and are observed in the data prior to 2012. This includes three measures: 'realized years per worker' aggregates the annual treatment effect for each cohort, as estimated in Table 2, and divides it by the number of workers in that cohort in 1991; 'Percent increase in worker-years' takes the worker-year as the level of analysis and estimates the lifetime treatment effect relative to all years worked by that cohort prior to 2012; and 'Realized years per person' distributes the lifetime treatment effect across the full cohort population in 1991. 'Potential' lifetime adjustment measures project annual adjustments past the available data as described in the text of this appendix. These estimates are comparable to the realized estimates calculated as 'years per person.'

Table A.4: Lifetime Adjustments in Labor-market Participation: Rural Men

	Age in 1991				
	25-34	35-44	45-54	55-59	60-69
Years per worker	0.23	1.15	2.33	1.62	0.45
Pct Increase in worker-years	0.96	4.11	8.64	7.32	2.57
Years per person	0.22	1.11	2.19	1.48	0.37
Upper bound (70), years per person	0.80	2.80	3.83	2.60	0.53

Notes. This table presents estimates of realized and potential lifetime labor supply adjustments for five-year cohorts of rural men. 'Realized' lifetime adjustments refer to increases in labor market participation that occurred and are observed in the data prior to 2012. This includes three measures: 'realized years per worker' aggregates the annual treatment effect for each cohort, as estimated in Table 2, and divides it by the number of workers in that cohort in 1991; 'Percent increase in worker-years' takes the worker-year as the level of analysis and estimates the lifetime treatment effect relative to all years worked by that cohort prior to 2012; and 'Realized years per person' distributes the lifetime treatment effect across the full cohort population in 1991. 'Potential' lifetime adjustment measures project annual adjustments past the available data as described in the text of this appendix. These estimates are comparable to the realized estimates calculated as 'years per person.'

Table A.5: Lifetime Adjustments in Labor-market Participation: Young rural Men

	Age in 1991				
	25-26	27-28	29-30	31-32	33-34
Realized adjustments					
Years per worker	0.32	0.12	0.10	0.26	0.45
Pct Increase in worker-years	1.65	0.56	0.44	1.07	1.78
Years per person	0.31	0.11	0.10	0.26	0.44
Potential adjustments					
Upper bound (70), years per person	1.17	0.56	0.48	1.01	1.21
Estimate given discontinuity					

Notes. This table presents estimates of realized and potential lifetime labor supply adjustments for two-year cohorts of rural men who were younger than 35 in 1991. 'Realized' lifetime adjustments refer to increases in labor market participation that occurred and are observed in the data prior to 2012. This includes three measures: 'realized years per worker' aggregates the annual treatment effect for each cohort, as estimated in Table 2, and divides it by the number of workers in that cohort in 1991; 'Percent increase in worker-years' takes the worker-year as the level of analysis and estimates the lifetime treatment effect relative to all years worked by that cohort prior to 2012; and 'Realized years per person' distributes the lifetime treatment effect across the full cohort population in 1991. 'Potential' lifetime adjustment measures project annual adjustments past the available data as described in the text of this appendix. These estimates are comparable to the realized estimates calculated as 'years per person.'

D Part-time Work

D.1 Rural married men

This appendix first provides further detail on labor-market participation patterns among rural married men. Figure A.11 is analagous to Figure 7 in the main text, but shows part-time work behavior among rural married men, rather than women. The left panel of Figure A.11 shows the fraction of the working population working part-time (between 1 and 29 hours per week) in urban and rural areas throughout our time-frame of interest. Visual inspection suggests an increase in part-time work among rural men following the pension expansion, and more formal difference-in-differences regressions in section D.2 confirm. The right panel shows the fraction of total population that works 45+ hours per week, 30-44 hours per week, and 1-29 hours per week. This figure suggests that, though there was a jump in the fraction of the population working 30-44 hours immediately following the constitutional reform in 1897 that was offset by a decrease in the fraction working 44+ hours per week, there is little evidence of a discontinuous change in the fraction of the population working in either full full-time category immediately after the implementation of the pension expansion in 1991. There was, however, a notable increase in part-time work in that year.

Figure A.11: Part-time Work Among Married Men

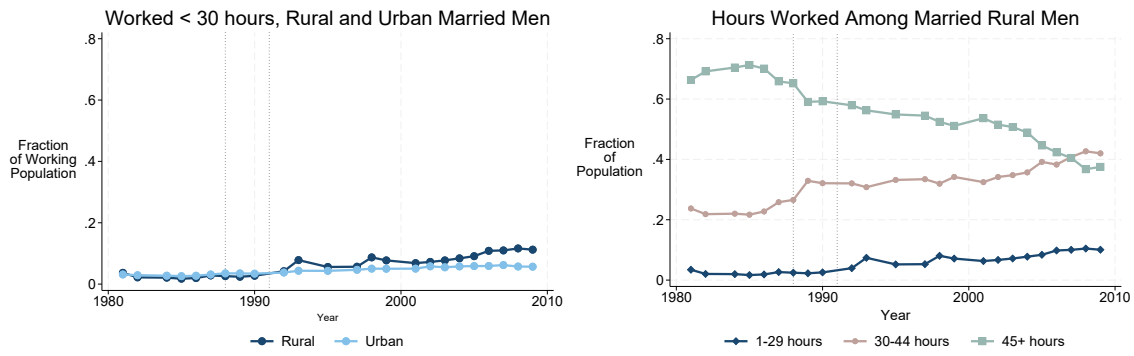

Notes. This figure describes trends in part-time work among married men who were aged 25-74 (first graph of each panel), younger than 55 in 1991 (second graph), and 55 and older in 1991 (third graph). Panel A shows the percent of the population that reported working in the reference week who worked less than 30 hours per week in rural (dark lines) and urban (light lines) areas. Panel B shows the percent of the rural married male population that reported working in the reference week who worked between 1 and 9 hours, 10 and 19 hours, 20 and 29 hours, 30 and 44 hours, and above 44 hours per week.

Column A in Figure A.12 shows histograms of part-time and full-time work among rural men in her three age-groups of interest: the always-age eligible (65-69), who were age-eligible to receive a pension before and after the 1991 reform; the newly-age eligible (60-65) who were not eligible to receive a pension prior to the reform; and the never-age eligible (55-59) who were not age-eligible to receive the pension either before or after the reform. This age decomposition mirrors the control groups used in (de Carvalho Filho, 2008), who explores the impact of this same pension reform on the labor supply of rural men. That paper identifies a negative wealth effect among rural men, using triple difference immediately before and after the reform, comparing men in rural versus urban. The third difference, based on age, nets out the effects of potential confounding trends

in rural area. The paper uses two different age groups as control groups: the never age-eligible (55-59) and the always age-eligible (65-69). Consistent with her findings, we see a large increase in the fraction of newly age-eligible men who were not participating in the labor market, and decrease in full-time work among those who remain in the market labor force. However, we also see notable changes in the labor-market patterns among men in the two comparison groups of interest. Never age-eligible men are less likely to work more than 45 hours a week, and more likely to work part-time. Always eligible men are similarly less likely to work more than 45 hours a week, but are also *more likely* to participate in the labor force at lower weekly hours. The triple-difference identification strategy relies on these differences in the control groups to be rural-specific time trends that are also affecting the newly-age eligible men.
While the patterns among the never age-eligible men are consistent with being a result of other reforms implemented with the 1988 constitution (see the part-time graph for men), and therefore necessary to extract from the newly-age eligible adjustments in labor supply, the patterns among the always age-eligible suggest that the age of men upon the reform impacted their experience of the reform. These findings support our conclusion that the pension expansion influenced labor supply decisions of those who were not directly included in age-eligibility as a result of the pension, perhaps through anticipatory or backward-looking responses.

A similar analysis for women suggests that the impacts are far smaller for the newly age-eligible, while the anticipatory labor supply responses of younger workers, and the ex-post labor supply responses of older workers, were far larger in comparison to men. This suggests that differences in the distribution of the utility weight on home production imply that the same reform has substantially different impacts on the labor market behavior of women versus men.

D.2 Empirical estimates of change in Part-time work

This section shows that rural workers were far more likely to work part-time following the reform than their urban counterparts. Rural women were 18 percentage points less likely to work after the reform than their urban counterparts, while rural men were seven percentage points less likely to work after the reform than their urban counterparts. For both men and women, this response was larger among workers who were older than their rural pension eligibility age.

In Tables A.6- A.8, we use the following specification to look at the extensive-margin labor supply response for people with shocks to own or spousal pension wealth:

$$y_{it} = \alpha \text{POST}_t + \beta \text{D}_i + \gamma \text{POST}_t \times \text{D}_i + \nu_t + \delta_s \tag{28}$$

POST is an indicator of whether individual i is observed before the reform (1981-1989) or after the reform (1991-2013). $Rural_i$ is a treatment indicator, equal to one if the individual lives in a rural area and zero otherwise. The outcome variable, y_{it} is an indicator equal to one if the individual worked less than 30 hours per week in the year t reference week, and the coefficient of interest is γ. The coefficients ν_t represent year fixed-effects, δ_s represent state fixed-effects. We limit the sample to individuals between ages 25 and 74 observed in year t who worked in the reference week, and run analysis separately for men and women.

Figure A.12: Three Age-Group Comparisons: Married Men and Women

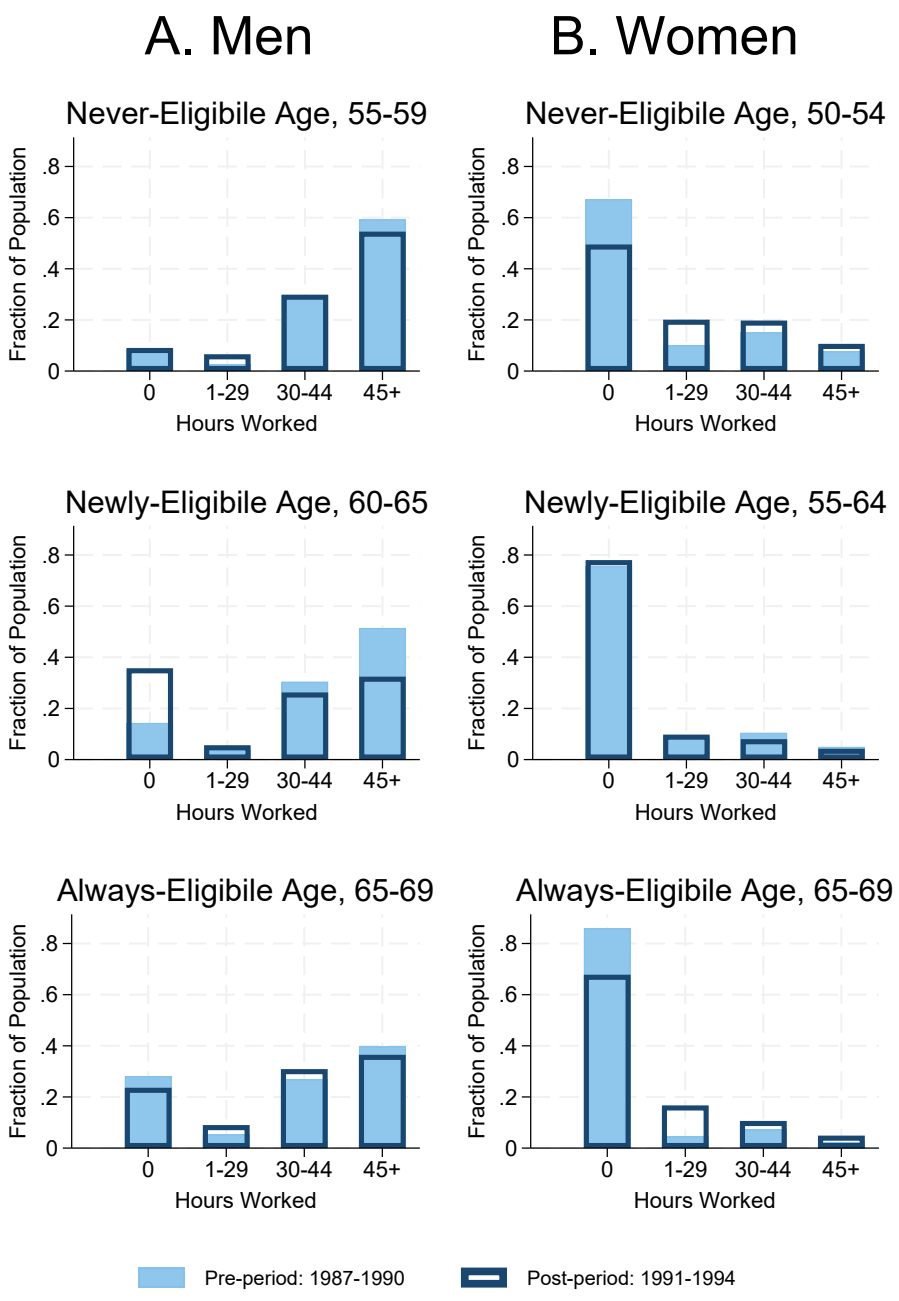

Notes. These figures describe labor-market participation for three age-groups of rural married men and women. The "never age-eligible" are too young to receive the pension, regardless of their work history, before and after the 1991 implementation of the expansion. The "newly-age eligible" are able to receive the pension following the reform, provided that they have worked the required number of years. The "never-age eligible" are old enough to receive the pension both before and after the reform, provided that they have worked the required number of years.

Figure A.13: Part-time work among rural and urban workers

Table A.6: Part-time work among Rural Workers

	Women				Men			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Rural*Post	0.19***	0.19***	0.18***	0.14***	0.07***	0.07***	0.07***	0.05***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Year FE		x	x	x		x	x	x
State FE			x	x			x	x
1981-1998 only				x				x
N	1648480	1648480	1648480	733253	2635977	2635977	2635977	1328845
R^2	0.033	0.047	0.048	0.030	0.018	0.030	0.031	0.020

Standard errors in parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Notes. POST is an indicator of whether individual i is observed before the reform (1981-1989) or after the reform (1991-2013). $Rural_i$ is a treatment indicator, equal to one if the individual lives in a rural area and zero otherwise. We limit the sample to individuals between ages 25 and 74 observed in year t who worked in the reference week. The sample includes both male and female workers.

Table A.7: Part-time work among Rural Women, Before and After Retirement

	Pre-retirment, < 55				Post-retirement, 55+			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Rural*Post	0.18***	0.18***	0.18***	0.13***	0.24***	0.24***	0.24***	0.14***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)
Year FE		x	x	x		x	x	x
State FE			x	x			x	x
1981-1998 only				x				x
N	1519399	1519399	1519399	687391	129081	129081	129081	45862
R^2	0.031	0.045	0.046	0.030	0.031	0.043	0.046	0.024

Standard errors in parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Notes. POST is an indicator of whether individual i is observed before the reform (1981-1989) or after the reform (1991-2013). $Rural_i$ is a treatment indicator, equal to one if the individual lives in a rural area and zero otherwise. We limit the sample to individuals between ages 25 and 74 observed in year t who worked in the reference week. The sample includes only female workers.

Table A.8: Part-time work among Rural Men, Before and After Retirement

	Pre-retirment, < 60				Post-retirement, 60+			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Rural*Post	0.07***	0.07***	0.07***	0.05***	0.09***	0.09***	0.09***	0.04***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Year FE		x	x	x		x	x	x
State FE			x	x			x	x
1981-1998 only				x				x
N	2478783	2478783	2478783	1257448	157194	157194	157194	71397
R^2	0.018	0.030	0.031	0.021	0.016	0.024	0.028	0.014

Standard errors in parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Notes. POST is an indicator of whether individual i is observed before the reform (1981-1989) or after the reform (1991-2013). $Rural_i$ is a treatment indicator, equal to one if the individual lives in a rural area and zero otherwise. We limit the sample to individuals between ages 25 and 74 observed in year t who worked in the reference week. The sample includes only male workers.

E Pensions and child-bearing

In this section, we show that rural female workers were less likely to have a child following the reform than their urban counterparts. We use the following specification to examine whether a rural female is less likely to have a child after the policy:

$$y_{it} = \alpha \text{POST}_t + \beta \text{Rural}_i + \gamma \text{POST}_t \times \text{Rural}_i + \nu_t + \delta_s \tag{29}$$

POST is an indicator of whether individual i is observed before the reform (1981-1989) or after the reform (1991-2013). $Rural_i$ is a treatment indicator, equal to one if the female lives in a rural area and zero otherwise. The outcome variable, y_{it} is an indicator equal to one if the female has a child in year t, and the coefficient of interest is γ. The coefficients ν_t represent year fixed-effects, δ_s represent state fixed-effects. We limit the sample to individuals between ages 25 and 74 observed in year t.

With the reform, rural women increases their labor supply, which in return would affect and reduce women's fertility behavior. In columns 1 and 2, we find negative difference-in-differences coefficients. When we examine women aged below 65 (column 3) and over 65 (column 4) separately, we find that the results are mainly driven by younger women. (Danzer & Zyska, 2022) similarly find that the reform reduces the propensity of childbearing of women.

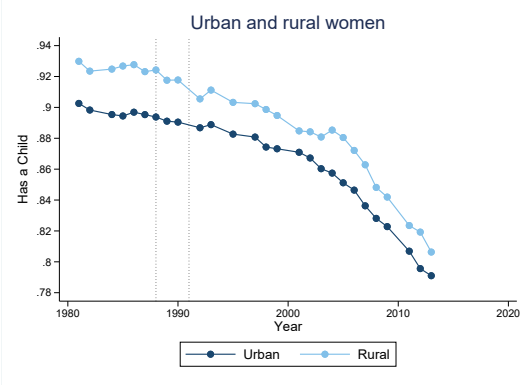

Figure A.14: Fertility among rural and urban women

Table A.9: Childbearing decision among women

	(1) Age No FE	(2) age and UF FE	(3) < 65	(4) >=65
Rural*Post	-0.0044***	-0.0043***	-0.0053***	0.0031
	(0.0008)	(0.0007)	(0.0007)	(0.0059)
N	5215323	5215048	5003507	211541
R^2	0.006	0.148	0.125	0.032

Standard errors in parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Notes. POST is an indicator of whether individual i is observed before the reform (1981-1989) or after the reform (1991-2013). $Rural_i$ is a treatment indicator, equal to one if the individual lives in a rural area and zero otherwise. We limit the sample to individuals between ages 25 and 74 observed in year t.