Who Earns? Universal Pre-K and the Division of Labor Within Households

Kosiso Madu[1]

October 2026

Access latest version here

Abstract

Universal pre-kindergarten (UPK) can increase maternal employment, but whether these employment gains change how earnings are divided within couples is less understood. This paper examines how UPK affects the division of market work and earnings within married couples. Using a stacked triple-difference design that exploits variation in the timing of UPK expansions across five U.S. jurisdictions and children's age eligibility, I find that UPK increases wives' share of couple earnings by 3.4 percentage points, or 11.3 percent relative to the pre-expansion mean. The increase is accompanied by higher labor force participation among wives and a larger share of couple market work, while husbands' labor force participation remains largely unchanged. I also find larger effects on wives' relative earnings below earnings parity. These findings show that the effects of public childcare extend beyond maternal employment to how married couples divide market work and earnings.

JEL classifications: J13, J16, J22.

Keywords: Universal pre-kindergarten, childcare, household labor supply, gender inequality.

[1]Department of Economics, University of Delaware. Email: kemadu@udel.edu.

1 Introduction

Parenthood changes the relative labor-market outcomes of women and men. Mothers reduce employment and work hours substantially more than fathers, and earnings gaps within couples widen after children arrive (Angelov, Johansson, and Lindahl, 2016; Kleven, Landais, and Søgaard, 2019; Goldin, Kerr, and Olivetti, 2022). Expanding access to public childcare and early education can increase maternal employment and labor supply (Gelbach, 2002; Cascio, 2009; Baker, Gruber, and Milligan, 2008; Humphries et al., 2024; Jackson, Turner, and Bastian, 2025). However, whether these employment gains change wives' relative contributions to couple earnings is less understood. A mother may enter employment while her husband continues to account for most of the couple's earnings. Similarly, spouses may change their relative earnings through adjustments in hours worked or other labor-market decisions without either spouse changing employment status.

The distribution of earnings within couples matters for reasons beyond whether both spouses earn income. Even among two-earner households, spouses can differ substantially in their contributions to household earnings. In the sample studied here, wives account for about 30 percent of couple earnings before the UPK expansions. Bargaining models of the household suggest that spouses' economic resources and opportunities may influence how household resources are allocated (Manser and Brown, 1980; McElroy and Horney, 1981; Lundberg and Pollak, 1993). Consistent with this, exogenous changes in which spouse receives income alter how the household spends it (Lundberg, Pollak, and Wales, 1997), and differences in spouses' relative earnings are associated with differences in consumption (Ahn and Koh, 2022). Against this background, this paper examines whether expanding access to universal pre-kindergarten (UPK) changes how married couples divide market work and earnings, including whether changes in wives' relative earnings extend beyond employment entry.

I study this question using the expansion of UPK across U.S. jurisdictions: Iowa, Wisconsin, the District of Columbia, Vermont, and New York City. Using data from the American Community Survey (ACS), I estimate a stacked triple-difference specification that exploits variation in the timing of UPK expansions and children's age eligibility. I compare households whose youngest child is eligible for UPK with households whose youngest child is aged six

to eight, before and after the expansion, in treated and never-treated jurisdictions. Because treatment occurs in only five jurisdictions, I supplement conventional clustered standard errors with randomization inference.

The main result shows that UPK increases wives' share of couple earnings by 3.4 percentage points, or 11.3 percent relative to the pre-expansion mean. The probability that the husband is the primary earner falls by 5.9 percentage points, while the probability that the wife is the primary earner rises by 3.4 percentage points. These changes are accompanied by an increase in wives' labor force participation and their share of couple market hours and weeks worked, while husbands' labor force participation remains largely unchanged. Taken together, these findings show that UPK shifts the division of earnings and market work toward wives.

I then examine whether the change in wives' earnings shares extends beyond the employment-entry margin. Restricting the sample to couples in which wives report positive earnings, I find a similar increase of 3.4 percentage points in wives' share of couple earnings. This suggests that the overall effect is not explained solely by women entering employment. I also examine how the response varies across the distribution of wives' earnings shares. UPK increases the probability that a wife earns more than 40 percent of couple earnings by 5.9 percentage points, but the effect falls to 3.0 percentage points at the 50 percent threshold. These results suggest that UPK shifts wives' relative earnings toward parity, with a smaller increase in the probability that wives earn at least as much as their husbands.

This paper makes two main contributions. First, it extends the literature on public childcare and maternal labor supply by examining how childcare access affects the distribution of earnings within couples (Gelbach, 2002; Cascio, 2009; Baker, Gruber, and Milligan, 2008; Humphries et al., 2024; Jackson, Turner, and Bastian, 2025). Existing studies have shown that public childcare can increase mothers' employment and earnings, but these changes do not necessarily tell us whether their contributions to couple earnings increase relative to their spouses'. I show that UPK increases wives' share of couple earnings, with a similar effect among couples in which wives report positive earnings. This finding suggests that the effects of childcare access on the distribution of earnings within couples are not explained solely by mothers entering employment.

Second, the paper contributes to the literature on the division of work within households. Previous research has documented how children affect specialization between spouses and the gender earnings gap (Juhn and McCue, 2017; Angelov, Johansson, and Lindahl, 2016; Kleven, Landais, and Søgaard, 2019). I show that expanding access to childcare changes how married couples divide market work, with wives accounting for a larger share of couple hours and weeks worked while husbands' labor force participation remains largely unchanged. I also find that the effects on wives' relative earnings are larger below earnings parity, connecting the findings to research on relative earnings within marriage (Bertrand, Kamenica, and Pan, 2015). These findings show that the division of market work and earnings within marriage responds to changes in childcare constraints, rather than remaining fixed when mothers increase their labor supply.

2 Background of UPK

Publicly funded early childhood education in the United States takes several forms. Head Start primarily serves children from economically disadvantaged families (U.S. Department of Health and Human Services, 2026). Many state-funded pre-kindergarten programs are also targeted, with eligibility based on family income, disability status, language needs, or other measures of disadvantage. In the 2024–2025 school year, 32 of the 63 state-funded preschool programs tracked by the National Institute for Early Education Research (NIEER) had an income requirement (Friedman-Krauss et al., 2026). These programs can reduce childcare costs for families that qualify, but they are not available to all families with preschool-aged children.

Universal pre-kindergarten (UPK) differs from targeted programs because eligibility is generally based on a child's age and residence rather than family income. This makes publicly funded pre-K available to a broader group of families, including those who would not qualify for targeted programs. However, universal eligibility does not necessarily mean universal enrollment or access to a family's preferred program. The number of available seats, whether local school districts participate, and whether programs operate part-day or full-day can affect how many eligible children attend. When demand exceeds available seats, programs

differ in how they allocate places and whether districts are required to expand capacity.[2] Enrollment therefore varies even across programs considered universal (Friedman-Krauss et al., 2026; Barnett and Gomez, 2016).

Publicly funded pre-K is not limited to public schools. Programs may also operate through private childcare centers, Head Start programs, and other community-based providers. All five jurisdictions in this study allow qualified private or community-based providers to offer publicly funded pre-K, although the arrangements differ across programs. Public funding can therefore increase access by creating new places in public schools, expanding existing programs, or funding places at participating private providers. In the latter case, existing private providers may serve children whose preschool education is now publicly funded.

UPK programs also differ in the number of hours they provide. Some operate for only part of the school day, while others provide care for most of the regular school day. For example, Vermont's Act 166 provides funding for at least ten hours of pre-kindergarten per week for 35 weeks annually, while New York City's 2014 expansion provided six hours and twenty minutes per day for 180 school days (Vermont General Assembly, 2014; City of New York, 2014). These differences matter for how much childcare families receive. A part-day program may reduce childcare costs without providing enough hours to cover a parent's work schedule, while a full-day program may reduce the need for additional care. Even full-day programs may not cover early mornings, evenings, school holidays, or other times when parents work.

The effect of UPK on childcare arrangements also depends on what families would have done without the program. Some children may otherwise have attended private preschool or another paid childcare arrangement. Others may have received care from a parent or relative. For families already paying for preschool, publicly funded pre-K may primarily

[2]In Iowa, participating school districts determine their enrollment procedures and are not required to open additional classrooms when demand exceeds available seats. In Wisconsin, districts may choose whether to offer four-year-old kindergarten, but participating districts are required to make the program available to eligible resident children. Districts operating programs in 2007–2008 that did not meet this requirement were given until 2013–2014 to comply. In the District of Columbia, oversubscribed schools use lotteries and waitlists, and children may receive offers as seats become available. In Vermont, school districts fund eligible children attending qualified programs but are not required to create additional seats when programs reach capacity. In New York City, admissions priorities determine placements, and children may receive offers at alternative programs or remain on waitlists for preferred locations. These arrangements do not necessarily imply that eligible children are unable to obtain a place anywhere in their jurisdiction.

reduce childcare expenses. For families providing care themselves, it may reduce the time parents need to spend caring for children during the school day. These differences mean that an increase in public pre-K enrollment does not necessarily represent an equal increase in childcare hours available to every family. They also help explain why access to the same program may lead to different labor-supply responses across households.

UPK has expanded across the United States at different times and through different approaches. I follow Jackson, Turner, and Bastian (2025) in identifying large expansions as programs open to children regardless of family income or other characteristics and enrolling at least 50 percent of four-year-olds by 2019. They identify nine state and local expansions: Georgia, Oklahoma, West Virginia, Florida, Iowa, Wisconsin, the District of Columbia, Vermont, and New York City. Some of these programs introduced publicly funded preschool to jurisdictions where few children previously attended, while others expanded existing programs by increasing district participation, available seats, or hours of care. The expansion years therefore capture substantial changes in access rather than necessarily the first year in which publicly funded pre-K was offered. Table 1 reports the expansion years, eligible ages, and enrollment rates for these programs. My analysis uses five of the expansions, as discussed in Section 4.

3 Empirical Strategy

I estimate the effect of UPK expansions using a stacked triple-difference design. The design combines variation in the timing of UPK expansions across jurisdictions with variation in the age of the youngest child. A difference-in-differences approach would compare households with UPK-eligible children in adopting and never-adopting jurisdictions before and after an expansion. However, changes in employment and earnings during this period may also reflect other developments in adopting jurisdictions. I therefore add households whose youngest child is aged six to eight as a comparison group. The triple-difference design compares changes between eligible-child and older-child households across adopting and never-adopting jurisdictions, allowing me to account for changes that affect both age groups within a jurisdiction (Olden and Møen, 2022).

I define eligibility using the age of the youngest child in the household. Mothers' employment tends to increase as their youngest child gets older (Lubotsky and Qureshi, 2018), and labor-supply responses to public policies also differ by children's ages (Michelmore and Pilkauskas, 2021). Using the age of the youngest child allows me to focus on households where UPK eligibility may reduce childcare constraints without the continued need to care for younger children. In Iowa, Wisconsin, and New York City, eligible households have a youngest child aged four; in the District of Columbia and Vermont, they have a youngest child aged three or four.

I use households whose youngest child is aged six to eight as the comparison group. These households have aged out of pre-K but remain relatively close in the childrearing life cycle to households with eligible children. I exclude households whose youngest child is five because these children may be in pre-K or kindergarten, depending on their birth month and the timing of the interview. Using older children also reduces the likelihood that the comparison group is responding to changes in preschool access at the same time as eligible households.

I estimate the following triple-difference specification:

$$Y_{ikstg} = \beta \left(\text{Post}_{stg} \times \text{Elig}_{kg} \right) + \alpha_{stg} + \theta_{ksg} + \delta_{ktg} + \varepsilon_{ikstg}, \tag{1}$$

where Y_{ikstg} is the outcome for household i, whose youngest child is age k, in jurisdiction s, year t, and stack g. Elig_{kg} equals one when age k is eligible for UPK in stack g. Post_{stg} equals one for the adopting jurisdiction in stack g after UPK expansion.

The specification includes jurisdiction-by-year fixed effects α_{stg}, child-age-by-year fixed effects δ_{ktg}, and child-age-by-jurisdiction fixed effects θ_{ksg}, all defined within each stack. Jurisdiction-by-year fixed effects absorb changes that affect all households in a jurisdiction in a given year, such as local labor-market conditions or other policies. Child-age-by-year fixed effects allow outcomes for households with children of different ages to change differently over time. Child-age-by-jurisdiction fixed effects allow the relationship between child age and the outcome to differ across jurisdictions. The coefficient β therefore measures the post-expansion change for households with UPK-eligible children relative to households with

older children in the same jurisdiction and to the corresponding difference in never-adopting jurisdictions (Olden and Møen, 2022).

To examine changes around the UPK expansion, I replace the post-treatment indicator with a set of event-time indicators:

$$Y_{ikstg} = \sum_{\substack{\ell=-4 \\ \ell \neq -1}}^{4} \beta_\ell \left(D^\ell_{stg} \times \text{Elig}_{kg} \right) + \alpha_{stg} + \theta_{ksg} + \delta_{ktg} + \varepsilon_{ikstg}, \tag{2}$$

where D^ℓ_{stg} equals one for the treated jurisdiction in stack g in event year ℓ and zero otherwise. Event year -1 is omitted, so each β_ℓ is measured relative to the year immediately before the UPK expansion.

The main identifying assumption is that, in the absence of UPK expansions, the difference in outcomes between eligible-child and older-child households would have followed similar trends in adopting and never-adopting jurisdictions. This does not require outcomes to follow the same trends across jurisdictions. Changes in local labor-market conditions, for example, may affect both groups within a jurisdiction. However, identification requires that no other changes at the time of the expansion differentially affect households with UPK-eligible children in adopting jurisdictions (Olden and Møen, 2022).

I examine this assumption using the pre-treatment event-study coefficients, which show whether the eligible-versus-older difference was changing differently in adopting and comparison jurisdictions before the expansion.

Identification also requires that households do not change their labor-market behavior before the expansion in anticipation of UPK access. The pre-treatment event-study estimates provide some evidence on possible anticipation, although they cannot rule out anticipatory responses that are too small to detect (Callaway and Sant'Anna, 2021; Roth et al., 2023).

A further concern is that UPK expansions could affect which households enter the sample. The analysis is restricted to married couples, and eligibility is determined by the age of the youngest child. If UPK affects marriage or fertility, the composition of the sample may

change after an expansion. I examine changes in marriage, fertility, and observable household characteristics following the expansions. I discuss these checks after the main results and report the estimates in the Appendix.

Because the five UPK expansions occur in different years, I estimate equations (1) and (2) using a stacked difference-in-differences design. For each expansion, I construct a separate stack containing the treated jurisdiction and the never-adopting comparison jurisdictions over the four years before and four years after the expansion. I then combine the five stacks and allow the fixed effects in equations (1) and (2) to vary by stack (Cengiz et al., 2019).

The stacked design compares each adopting jurisdiction with never-adopting jurisdictions over the same event-time window and avoids using already-treated jurisdictions as controls. Such comparisons can arise in conventional two-way fixed-effects models with staggered treatment timing and can be difficult to interpret when treatment effects differ across groups or over time (Goodman-Bacon, 2021; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021). The separate stacks also allow eligibility to reflect the ages covered by each program. I report conventional two-way fixed-effects estimates as a robustness check in the Appendix.

One concern with the triple-difference design is that some households in the older-child comparison group may have been exposed to UPK in an earlier year. For example, a household whose youngest child is six two years after an expansion may have had a four-year-old when the program began. If the effects of UPK continue after the child leaves pre-K, previous exposure among the comparison group could affect the estimated difference between currently eligible and older-child households.

I therefore also estimate an across-jurisdiction difference-in-differences specification that restricts the sample to households whose youngest child is age-eligible for UPK. I compare these households across adopting and never-adopting jurisdictions before and after the expansion.[3] The across-jurisdiction specification avoids relying on the older-child comparison

[3]The alternative specification is

$$Y_{istg} = \beta \text{Post}_{stg} + \alpha_{sg} + \gamma_{tg} + \varepsilon_{istg},$$

where α_{sg} and γ_{tg} are jurisdiction and year fixed effects within each stack. This specification does not use aged-out households as a comparison group. Results are reported in the Appendix.

group, but requires that outcomes for eligible-child households would have followed similar trends in adopting and never-adopting jurisdictions without UPK. Unlike the triple-difference specification, it cannot absorb jurisdiction-by-year specific shocks. I therefore use the triple difference as the preferred specification and report the across-jurisdiction estimates in the Appendix.

All regressions are weighted using ACS household weights, and standard errors are clustered at the jurisdiction level, where UPK treatment varies (Bertrand, Duflo, and Mullainathan, 2004; Cameron and Miller, 2015).

Because UPK treatment occurs in only five jurisdictions, conventional cluster-robust standard errors may be unreliable (Olden and Møen, 2022; Conley and Taber, 2011; Cameron and Miller, 2015; Ferman and Pinto, 2019). I therefore report jurisdiction-clustered standard errors alongside studentized randomization-inference p-values (MacKinnon and Webb, 2020). [4]

4 Data

The main data come from the American Community Survey (ACS). The analysis uses repeated cross-sectional data on children's ages, labor-force status, usual hours worked, and annual wage and salary income for both spouses.[5] All household-level estimates are weighted using ACS household weights.

[4]For the randomization-inference procedure, I keep the five observed treatment years and the structure of the stacked sample fixed. In each placebo draw, I select five distinct jurisdictions from the never-treated comparison pool and assign one of the observed treatment events to each jurisdiction. I then re-estimate the same stacked triple-difference specification using these placebo assignments. Repeating this procedure $B = 500$ times produces the placebo distribution of the estimated treatment effect. For each outcome, I calculate a two-sided randomization-inference p-value using the studentized test statistic:

$$p_{\text{RI}} = \frac{1 + \sum_{b=1}^{B} \mathbf{1}\{|t^b_{\text{placebo}}| \geq |t_{\text{actual}}|\}}{B+1},$$

where each t-statistic is the estimated coefficient divided by its jurisdiction-clustered standard error. Randomization inference based on the coefficient itself gives similar, albeit less conservative, results.

[5]Annual wage and salary income is converted to 2010 dollars. Earnings refer to the previous twelve months, while hours are usual hours worked per week. The ACS is a repeated cross-section, so households are observed once and cannot be followed as children age out of eligibility.

I also use data from the National Institute for Early Education Research (NIEER) to describe state-funded pre-kindergarten programs. The NIEER data provide information on program eligibility and the share of children enrolled in state-funded pre-K. I use these data to document the reach of the UPK programs, construct the enrollment rates reported in Table 1, and examine changes in enrollment following the expansions.

The main sample consists of households containing one opposite-sex married couple with children.[6] I restrict the sample based on the age of the youngest child because UPK eligibility is determined by child age. In Iowa, Wisconsin, and New York City, households are age-eligible when the youngest child is four. In the District of Columbia and Vermont, households are age-eligible when the youngest child is three or four. The comparison group for the triple-difference design consists of households whose youngest child is between six and eight. Households whose youngest child is five are excluded.

I follow Jackson, Turner, and Bastian (2025) in identifying large UPK expansions. Of the nine expansions they identify, five enter my analysis: Iowa (2007), Wisconsin (2008), the District of Columbia (2009), Vermont (2014), and New York City (2014). Georgia and Oklahoma expanded before the available ACS data, while West Virginia's 2002 expansion leaves insufficient pre-expansion observations. Florida expanded in 2005, so its entire pre-expansion period would come from the smaller ACS samples available before 2005.

For each treated jurisdiction, I use observations from four years before through four years after the UPK expansion, as described in Section 3.[7] The resulting sample spans 2003–2018. For the preferred triple-difference specification, the sample includes both age-eligible households and households with a youngest child aged six to eight. The double-difference specification uses only households whose youngest child is four in all jurisdictions.

The main outcome is the wife's share of couple earnings. I define this as the wife's annual wage and salary income divided by the combined wage and salary income of both spouses. The measure ranges from zero to one and is defined for households with positive total couple

[6] I restrict wives to ages 18–55 and husbands to ages 18–65 and exclude couples with an age difference of more than 25 years.

[7] New York City is treated separately from the rest of New York State and is identified using the five borough counties.

earnings. A higher value means that the wife accounts for a larger share of the couple's earnings.

I use labor-force participation to examine changes on the extensive margin. Wife labor-force participation equals one if the wife is in the labor force and zero otherwise. I define the same measure separately for husbands. To examine the effects of UPK on a broader sample of women with age-eligible children, regardless of marital status, I also use an employment indicator, which equals one if the woman is employed and zero if she is unemployed or out of the labor force. The employment indicator is defined for all women rather than only labor-force participants.

Allocation of market time is measured in two ways. The wife's share of couple weeks is her weeks worked divided by the total weeks worked by both spouses. Weeks are reported in intervals in the ACS, and I assign each interval its midpoint.[8] The wife's share of couple estimated annual hours is her usual weekly hours multiplied by her weeks, divided by the same product summed across both spouses. Annual hours reflect adjustment on both the weekly and the yearly margin, while weeks reflect only the latter. Each measure is defined for couples with a positive total, and a higher value means that the wife accounts for a larger share of the couple's market work.

Several additional outcomes are used to describe changes in the allocation of work and earnings within the household. Husband primary earner equals one when the husband receives at least 60 percent of couple earnings, while wife primary earner equals one when the wife receives at least 60 percent. When either spouse's share of couple earnings is strictly between 40 and 60 percent, both indicators equal zero. I also report the wife's real annual wage and salary income in levels.

For the first stage, I use data from NIEER and the ACS to examine changes in preschool enrollment following UPK expansions. The NIEER data measure the share of four-year-olds enrolled in state-funded pre-K, while the ACS provides three measures of preschool

[8]The ACS reports weeks worked in the previous year in seven categories. I replace each with the midpoint of its range: 0 weeks with 0, 1–13 weeks with 7, 14–26 with 20, 27–39 with 33, 40–47 with 43.5, 48–49 with 48.5, and 50–52 with 51.

attendance among eligible children: attendance at any preschool, public preschool, and private preschool.[9]

Table 2 compares households with age-eligible children in treated and never-treated jurisdictions before the UPK expansions. The two groups have similar labor-market outcomes. Wives' labor-force participation is 64.9 percent in treated jurisdictions and 63.7 percent in never-treated jurisdictions. The wife's share of couple earnings is 0.299 and 0.287, respectively, while her share of couple annual hours is 0.306 and 0.296. The probability that the husband is the primary earner is 65.0 and 65.6 percent. Appendix Table A1 reports the corresponding statistics for households whose youngest child is aged six to eight, along with household characteristics and the differences between the two age-group gaps.

5 Results

5.1 First Stage

I first examine whether the UPK expansions increased preschool enrollment. Table 3 reports estimates using two data sources. Column 1 uses NIEER data on the share of four-year-olds enrolled in state-funded pre-kindergarten. Columns 2–4 use ACS data on attendance at any preschool, public preschool, and private preschool.

Both use a double difference rather than the triple difference used elsewhere in the paper. The NIEER panel is a state-year series with no child-age dimension. In the ACS columns, the comparison group of children aged six to eight is mostly in elementary school, so the age comparison carries no variation.

Column 1 shows that state-funded pre-K enrollment increased by 17.9 percentage points, relative to a pre-expansion mean of 38.2 percent. Figure 1 shows how enrollment changed around the expansions. The estimates are small before expansion, while enrollment increases gradually afterward. The estimated increase is 2 percentage points in the expansion year, 13

[9]The ACS distinguishes public and private schools based on how they are operated and funded, rather than who pays for an individual child's place (Census Bureau, 2010). Because UPK may fund places at private providers, children attending publicly funded pre-K may still be reported as attending private preschool.

percentage points one year later, and 27 percentage points by the fourth year. This pattern is consistent with programs expanding capacity over several years.

The NIEER analysis includes Iowa, Wisconsin, the District of Columbia, and Vermont. New York State is excluded entirely because NIEER does not report enrollment separately for New York City, making it impossible to distinguish the city's expansion from enrollment elsewhere in the state. The ACS analysis includes all five treated jurisdictions because New York City can be identified using county codes.

Columns 2–4 report the ACS estimates using a double-difference specification for four-year-old children. Public preschool attendance increases by 3.1 percentage points, while attendance at any preschool increases by 2.6 percentage points, although the latter estimate is not statistically significant. Private preschool attendance falls by an estimated 0.5 percentage points, but the change is also not statistically significant.[10] These results provide evidence of increased public preschool attendance, but do not establish whether the expansion displaced enrollment in private preschool. Appendix Figure A1 reports the event-study estimates for public preschool attendance.

The difference between the NIEER and ACS estimates can be attributed to what each source measures. NIEER includes children enrolled in state-funded programs regardless of whether the provider is public or private. The ACS instead records whether a child attends a public or private preschool. Publicly funded places offered by private providers may therefore be reported as private preschool attendance in the ACS.

Before the expansions, 68 percent of four-year-olds in the estimation sample already attended preschool. The estimated increase in total preschool attendance is smaller than the increase in state-funded enrollment, suggesting that some of the expansion may have benefited children who would otherwise have attended preschool without public funding. UPK may therefore have reduced childcare costs for some families in addition to increasing access.

I do not scale the household estimates by the first stage because NIEER measures state-funded enrollment for a broader population, while the ACS does not distinguish publicly

[10]Because public and private preschool attendance partition total preschool attendance, the estimated effects add up: $0.0311 - 0.0049 = 0.0262$.

funded UPK places from privately funded preschool attendance. Because treatment is based on eligibility rather than actual enrollment, the household estimates are intention-to-treat effects.

5.2 Women's Labor Supply

Before examining the division of work and earnings within married couples, I estimate the effect of UPK on women's labor supply more broadly. Appendix Table A2 reports estimates from the stacked triple-difference specification using person-level ACS data for women aged 18–55 whose youngest child is age-eligible or in the comparison group, regardless of marital status. UPK increases labor-force participation by 3.9 percentage points, or 5.6 percent relative to the pre-expansion mean, and employment by 3.1 percentage points. These findings are consistent with previous evidence that expanded access to public pre-kindergarten increases maternal employment (Jackson, Turner, and Bastian, 2025).

I then examine estimates by marital status. Among married women, UPK increases labor-force participation by 5.8 percentage points, compared with 3.4 percentage points among cohabiting women and 0.9 percentage points among single women. The corresponding employment estimates are 5.3, 3.8, and -1.3 percentage points, respectively, but the last two are not distinguishable from zero.

For married women who also meet the couple-sample restrictions, the estimated increase in labor-force participation is 5.6 percentage points, close to the 5.5 percentage points obtained from the household-level sample. I focus on married couples in the remainder of the analysis to examine how changes in women's labor supply are reflected in the division of earnings and market work between spouses.

5.3 Division of Earnings and Market Work Within Couples

Before reporting average treatment effects, I examine the event-study estimates in Figures 2 and 3. The pre-expansion estimates provide a check on whether the difference between age-eligible and comparison households was already changing differently in treated and never-treated jurisdictions. The estimates show no clear pattern of differential trends before the

UPK expansions. Following the expansions, the wife's share of couple earnings increases, while the probability that the husband is the primary earner declines. Wives' labor-force participation and their share of couple annual hours also increase, while husbands' participation shows no significant change.

Table 4 reports the average treatment effects. UPK increases the wife's share of couple earnings by 3.4 percentage points, relative to the pre-expansion mean of 29.9 percent. The change is also evident in measures of who earns the household's income. The probability that the husband earns at least 60 percent of couple earnings falls by 5.9 percentage points from a pre-expansion mean of 65.0 percent, a decline of approximately 9 percent. The probability that the wife earns at least 60 percent rises by 3.4 percentage points from a pre-expansion mean of 14.7 percent, an increase of approximately 23 percent. These results show a shift in the distribution of earnings toward wives, including changes in which spouse is the primary earner. These findings complement Humphries et al. (2024), who find that extended-day pre-kindergarten in New Haven increases earnings for both mothers and fathers and cannot reject equal effects across the two groups. While they use admissions lotteries to study individual parents' earnings, I examine how UPK expansions affect the division of earnings within married couples.

The allocation of market work changes in the same direction. The wife's share of couple weeks worked increases by 2.4 percentage points, while her share of estimated annual hours increases by 2.7 percentage points.

Table 5 reports the labor-force participation estimates. Wife labor-force participation increases by 5.5 percentage points, relative to a pre-expansion mean of 64.9 percent. My estimate is larger than the 1.1 percentage point increase in labor-force participation reported by Jackson, Turner, and Bastian (2025) for mothers of children under six relative to prime-age men. One difference is that I focus on married couples whose youngest child is age-eligible for UPK, whereas their sample includes mothers of young children more broadly. The estimated effect on husbands' labor-force participation is close to zero and statistically insignificant, consistent with their finding of little change among prime-age men. The increase in wives' participation therefore does not appear to be accompanied by husbands leaving the labor

force.

Relative to their pre-expansion means, wives' labor-force participation rises by 8.5 percent, their share of couple weeks worked by 6.8 percent, their share of couple annual hours by 8.9 percent, and their share of couple earnings by 11.3 percent. These changes suggest that the increase in wives' relative earnings is accompanied by changes in both labor-force participation and the allocation of market work within couples.

I also examine wives' annual earnings in levels. The estimated effect is \$5,487 with a clustered standard error of \$2,448, but the studentized randomization inference p-value is 0.140. The estimate is therefore not statistically significant under randomization inference, unlike the main earnings-share estimates.[11] The evidence for an increase in wives' earnings levels is therefore less conclusive than the evidence for a change in their share of couple earnings.

Taken together, the results show that UPK increases wives' labor-force participation and their share of couple earnings and market work, without a detectable decline in husbands' labor-force participation.

5.4 Where in the Earnings Distribution Do the Effects Occur?

The estimated increase in wives' average share of couple earnings does not show where in the distribution the changes occur. To examine this, I estimate the preferred specification using indicators for whether the wife's earnings share exceeds different thresholds, ranging from zero to 0.90. The results are reported in Table 6 and Figure 4.

The estimated effects increase across the lower thresholds and reach their largest value at 40 percent. UPK increases the probability that a wife earns more than 40 percent of couple earnings by 5.9 percentage points, compared with 3.0 percentage points at the 50 percent threshold. The difference implies an increase of approximately 3.0 percentage points in the probability that the wife earns between 40 and 50 percent of couple earnings. The results therefore suggest that the increase in wives' relative earnings is particularly pronounced

[11] The earnings-level estimates are sensitive to observations at the top of the distribution. Excluding the top one percent reduces the estimated effect on wives' earnings to \$2,916 (s.e. \$1,339), while excluding the top five percent reduces it to \$1,046 (s.e. \$820).

below earnings parity, although there is also an increase in the probability of wives earning more than half.

This pattern is relevant to the literature on the division of earnings within couples. Bertrand, Kamenica, and Pan (2015) document a sharp decline in the density of wives' earnings shares above 50 percent and interpret it as evidence of gender identity norms. Binder and Lam (2022) challenge this interpretation, showing that couples reporting identical earnings can account for the discontinuity.

The ACS data also exhibit a mass point at equal earnings: 2.3 percent of couple observations report an earnings share of exactly 0.50. To examine whether this affects the threshold estimates, I compare indicators defined using strict and weak inequalities. The estimated effect at the 50 percent threshold is 0.0296 when defined as greater than 0.50 and 0.0292 when defined as greater than or equal to 0.50. Thus, the smaller estimated effect at parity is not sensitive to the inclusion of couples reporting identical earnings.

6 Sensitivity and Robustness

I conduct several checks to examine how sensitive the main findings are to alternative specifications, sample composition, and the choice of comparison jurisdictions.

I first examine whether the results depend on the preferred stacked triple-difference specification. Appendix Table A3 compares the preferred specification with three alternatives: a stacked double difference, a conventional two-way fixed-effects triple difference, and a two-way fixed-effects double difference.

The estimated effects are generally similar across the four specifications. The estimated effect on the wife's share of couple earnings ranges from 2.7 to 3.4 percentage points, while the effect on wife labor-force participation ranges from 4.8 to 5.5 percentage points. The wife's share of annual hours increases by 2.4 to 2.7 percentage points, husband-primary-earner status falls by 5.1 to 5.9 percentage points, and wife-primary-earner status rises by 2.3 to 3.8 percentage points. The wife's share of weeks worked also increases across all four specifications. Appendix Figure A2 reports event-study estimates from the across-

jurisdiction DiD specification for comparison with the preferred triple-difference design.

I next examine whether the results are driven by any single UPK expansion. Appendix Table A4 reports estimates from the preferred stacked triple-difference specification, excluding one treated jurisdiction and its corresponding stack at a time.

The estimates remain similar in sign and magnitude across all five exclusions. The effect on the wife's share of couple earnings ranges from 3.1 to 3.9 percentage points, compared with 3.4 percentage points in the full sample. The estimated increase in wives' labor-force participation ranges from 4.4 to 6.6 percentage points, while the decline in husband-primary-earner status ranges from 5.3 to 6.8 percentage points. The estimates for wives' share of market work also remain positive, while husbands' labor-force participation remains close to zero.

New York City and Wisconsin account for the largest shares of treated households, at 33.4 and 35.0 percent, respectively. Excluding either jurisdiction does not substantially change the findings, suggesting that the results are not driven by any single expansion.

Because the main analysis is restricted to married couples whose youngest child falls within particular age groups, a concern is that UPK expansions could affect who enters the estimation sample. Changes in marriage or fertility could alter the composition of eligible households, potentially affecting the estimated division of earnings and market work.

Appendix Table A5 examines this possibility using a person-level sample of women, regardless of marital status. The estimated effect on being married with a spouse present is 1.1 percentage points (s.e. 1.1), with little evidence of changes in cohabitation. I also examine fertility-related selection. In the age-restricted sample, the estimated effect on the number of children is -0.026 (s.e. 0.015). Because women who have another child may leave the age-eligible group, I repeat the analysis using a broader sample of women whose youngest child is aged zero to eight. The estimated effects on having a youngest child aged four and on the number of children are small and statistically insignificant. Overall, these estimates provide limited evidence that UPK substantially changes the composition of the estimation sample through marriage or fertility.

I also assess how different couples entering the estimation sample would need to be for changes in marriage to account for the main results. Appendix Table A6 reports the mean outcomes required among entering couples under a simple sample-composition exercise.[12] Using the estimated change in marriage, entrants would constitute approximately 1.7 percent of the post-expansion married sample under this assumption. I also consider a larger share of 5.0 percent, calculated using the upper end of the conventional 95 percent confidence interval.

Even under the larger assumed share, explaining the estimated effects on wives' labor-force participation and husband-primary-earner status would require entrant means outside the feasible range. Explaining the increase in wives' earnings share would require entrants to have a mean earnings share of 97.1 percent, compared with 29.9 percent among existing couples. These calculations suggest that changes in marriage are unlikely to explain the main findings through sample composition alone, even when using the larger assumed share of entrants.

I next examine whether changes in observable household characteristics account for the main findings. Panel A of Appendix Table A7 shows that both spouses are approximately four months older following UPK expansion, while the share of wives with a college degree increases by 2.5 percentage points. Panel B adds age, education, and race controls for both spouses to the preferred specification. The estimated effect on wives' earnings share changes from 3.37 to 3.19 percentage points, while the estimates for primary-earner status, labor-force participation, and wives' share of market work remain similar. Thus, although the composition of married couples changes along some observable dimensions, controlling for these characteristics does not substantially alter the main findings.

The comparison group includes jurisdictions with state-funded pre-kindergarten programs, some of which expanded enrollment during the study period. If these expansions also in-

[12]The calculation treats the post-expansion married sample as a mixture of existing and entering couples, following the logic of contamination bounds (Horowitz and Manski, 1995). If existing couples have mean outcome μ, entrants have mean m, and entrants account for share s of the post-expansion sample, the change attributable to composition alone is $s(m-\mu)$. Setting this equal to the estimated effect β gives the entrant mean required to explain it, $m = \mu + \beta/s$. The entrant share is approximated using the estimated change in marriage from Table A5, assuming entry without offsetting exits. Because the marriage estimate comes from a separate person-level sample, the implied entrant share is only a proxy for entry into the couple estimation sample. The exercise is therefore illustrative rather than a formal bound on the treatment effect.

creased access to preschool, they could reduce the estimated difference between treated and comparison jurisdictions.

To examine this concern, Appendix Table A8 progressively restricts the comparison group based on state-funded pre-K enrollment among four-year-olds. A jurisdiction is retained in a given stack only if its enrollment rate remains at or below the specified threshold throughout the nine-year window. The restrictions are applied separately within each stack, while the five treated jurisdictions remain unchanged.

The estimates are generally similar or somewhat larger as the comparison group becomes more restrictive. The estimated effect on wives' share of couple earnings increases from 3.37 percentage points in the baseline to 4.10 percentage points at the 5 percent threshold. The effect on wives' share of annual hours increases from 2.73 to 3.29 percentage points, while the effect on labor-force participation changes from 5.49 to 5.89 percentage points. The estimates for primary-earner status also remain similar across restrictions.

These patterns are consistent with the possibility that expanding pre-K access in comparison jurisdictions attenuates the baseline estimates. However, restricting the comparison group also reduces the number of jurisdictions and changes its composition. The results therefore do not establish that contamination explains the differences in estimates. Overall, the main findings are not sensitive to excluding jurisdictions with higher levels of state-funded pre-K enrollment.

7 Conclusion

This paper examines whether expanding access to universal pre-kindergarten changes how married couples divide market work and earnings. Using a stacked triple-difference design across five U.S. jurisdictions, I find that UPK increases wives' share of couple earnings by 3.4 percentage points, or 11.3 percent relative to the pre-expansion mean. The probability that the husband is the primary earner falls by 5.9 percentage points, while the probability that the wife is the primary earner rises by 3.4 percentage points. Wives' labor-force participation increases by 5.5 percentage points, while husbands' participation remains largely unchanged.

The increase in wives' earnings share is similar among couples in which wives report positive earnings, suggesting that the effect is not explained solely by entry into employment.

The effects on wives' relative earnings are particularly large below earnings parity. UPK increases the probability that a wife earns more than 40 percent of couple earnings by 5.9 percentage points, compared with 3.0 percentage points at the 50 percent threshold. The difference implies an increase of approximately 3.0 percentage points in the probability that wives earn more than 40 percent but no more than 50 percent of couple earnings. This pattern is relevant to Bertrand, Kamenica, and Pan (2015), who document a discontinuity in the distribution of wives' relative earnings around parity. Although the estimated response is larger below parity, the analysis does not establish why it differs across thresholds.

These findings relate to the literature on parenthood and household specialization. Angelov, Johansson, and Lindahl (2016) document how earnings differences between spouses widen following childbirth, while Juhn and McCue (2017) examine the persistence of gender earnings gaps associated with children. My findings show that the division of market work and earnings among couples with young children also responds to childcare policy. In contrast, Canaan (2022) finds that an expansion of extended paid parental leave increased household specialization, with mothers reducing employment and fathers increasing their work hours. Together, these findings suggest that family policies do not have a uniform effect on the division of market work among couples.

The results also have implications for how childcare policies are evaluated. Studies that focus on maternal employment or earnings alone cannot tell us whether spouses' relative contributions to household earnings have changed. My analysis shows that UPK affects both women's individual labor-market outcomes and the division of earnings and market work within couples. Examining these outcomes together provides a more complete picture of how families respond to expanded childcare access.

The analysis observes market work and earnings, not how spouses divide unpaid work within the home. It also captures responses around the UPK expansions rather than the longer-run division of work as children grow older. Whether the changes in market work are accompanied by changes in home production, and whether they persist beyond UPK eligibility, remain

questions for future research.

References

Ahn, So Yoon, and Yu Kyung Koh. 2022. "Spousal Bargaining Power and Consumption of Married Couples in the US: Evidence from Scanner Data." HCEO Working Paper No. 2022-010. Human Capital and Economic Opportunity Global Working Group.

Angelov, Nikolay, Per Johansson, and Erica Lindahl. 2016. "Parenthood and the Gender Gap in Pay." *Journal of Labor Economics* 34(3): 545–579. doi:10.1086/684851.

Baker, Michael, Jonathan Gruber, and Kevin Milligan. 2008. "Universal Child Care, Maternal Labor Supply, and Family Well-Being." *Journal of Political Economy* 116(4): 709–745. doi:10.1086/591908.

Barnett, W. Steven, and Rebecca Gomez. 2016. "Universal Pre-K: What Does It Mean and Who Provides It?" National Institute for Early Education Research, Rutgers University.

Barnett, W. Steven, Megan E. Carolan, Jen Fitzgerald, and James H. Squires. 2012. *The State of Preschool 2012: State Preschool Yearbook*. New Brunswick, NJ: National Institute for Early Education Research.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics* 119(1): 249–275.

Bertrand, Marianne, Emir Kamenica, and Jessica Pan. 2015. "Gender Identity and Relative Income within Households." *Quarterly Journal of Economics* 130(2): 571–614. doi:10.1093/qje/qjv001.

Binder, Ariel J., and David Lam. 2022. "Is There a Male-Breadwinner Norm? The Hazards of Inferring Preferences from Marriage Market Outcomes." *Journal of Human Resources*.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225(2): 200–230.

Cameron, A. Colin, and Douglas L. Miller. 2015. "A Practitioner's Guide to Cluster-Robust Inference." *Journal of Human Resources* 50(2): 317–372. doi:10.3368/jhr.50.2.317.

Canaan, Serena. 2022. "Parental Leave, Household Specialization and Children's Well-Being." *Labour Economics* 75: 102127. doi:10.1016/j.labeco.2022.102127.

Cascio, Elizabeth U. 2009. "Maternal Labor Supply and the Introduction of Kindergartens into American Public Schools." *Journal of Human Resources* 44(1): 140–170. doi:10.3368/jhr.44.1.140.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *Quarterly Journal of Economics* 134(3): 1405–1454.

Conley, Timothy G., and Christopher R. Taber. 2011. "Inference with 'Difference in Differences'

with a Small Number of Policy Changes." *Review of Economics and Statistics* 93(1): 113–125. doi:10.1162/REST_a_00049.

Council of the District of Columbia. 2008. "Pre-K Enhancement and Expansion Amendment Act of 2008." D.C. Law 17-202, effective July 18, 2008.

Ferman, Bruno, and Cristine Pinto. 2019. "Inference in Differences-in-Differences with Few Treated Groups and Heteroskedasticity." *Review of Economics and Statistics* 101(3): 452–467. doi:10.1162/rest_a_00759.

Friedman-Krauss, Allison H., W. Steven Barnett, Katherine S. Hodges, Karin A. Garver, G. G. Weisenfeld, Jennifer Duer, and Jessica Siegel. 2026. *The State of Preschool 2025: State Preschool Yearbook*. New Brunswick, NJ: National Institute for Early Education Research, Rutgers University.

Gelbach, Jonah B. 2002. "Public Schooling for Young Children and Maternal Labor Supply." *American Economic Review* 92(1): 307–322. doi:10.1257/000282802760015748.

Goldin, Claudia, Sari Pekkala Kerr, and Claudia Olivetti. 2022. "When the Kids Grow Up: Women's Employment and Earnings across the Family Cycle." NBER Working Paper No. 30323. National Bureau of Economic Research. doi:10.3386/w30323.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics* 225(2): 254–277.

Horowitz, Joel L., and Charles F. Manski. 1995. "Identification and Robustness with Contaminated and Corrupted Data." *Econometrica* 63(2): 281–302. doi:10.2307/2951627.

Humphries, John Eric, Christopher Neilson, Xiaoyang Ye, and Seth D. Zimmerman. 2024. "Parents' Earnings and the Returns to Universal Pre-Kindergarten." NBER Working Paper No. 33038. National Bureau of Economic Research. doi:10.3386/w33038.

Jackson, C. Kirabo, Julia A. Turner, and Jacob Bastian. 2025. "Universal Pre-K as Economic Stimulus: Evidence from Nine States and Large Cities in the U.S." NBER Working Paper No. 33767, revised April 2026. National Bureau of Economic Research. doi:10.3386/w33767.

Juhn, Chinhui, and Kristin McCue. 2017. "Specialization Then and Now: Marriage, Children, and the Gender Earnings Gap across Cohorts." *Journal of Economic Perspectives* 31(1): 183–204. doi:10.1257/jep.31.1.183.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard. 2019. "Children and Gender Inequality: Evidence from Denmark." *American Economic Journal: Applied Economics* 11(4): 181–209. doi:10.1257/app.20180010.

Lubotsky, Darren, and Javaeria A. Qureshi. 2018. "Assessing the Smooth Rise in Mothers' Employment as Children Age." *Journal of Human Capital* 12(4): 604–639. doi:10.1086/700077.

Lundberg, Shelly J., Robert A. Pollak, and Terence J. Wales. 1997. "Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit." *Journal of Human Resources* 32(3): 463–480. doi:10.2307/146179.

Lundberg, Shelly, and Robert A. Pollak. 1993. "Separate Spheres Bargaining and the Marriage Market." *Journal of Political Economy* 101(6): 988–1010. doi:10.1086/261912.

MacKinnon, James G., and Matthew D. Webb. 2020. "Randomization Inference for Difference-in-Differences with Few Treated Clusters." *Journal of Econometrics* 218(2): 435–450. doi:10.1016/j.jeconom.2020.04.024.

Manser, Marilyn, and Murray Brown. 1980. "Marriage and Household Decision-Making: A Bargaining Analysis." *International Economic Review* 21(1): 31–44. doi:10.2307/2526238.

McElroy, Marjorie B., and Mary Jean Horney. 1981. "Nash-Bargained Household Decisions: Toward a Generalization of the Theory of Demand." *International Economic Review* 22(2): 333–349.

Michelmore, Katherine, and Natasha Pilkauskas. 2021. "Tots and Teens: How Does Child's Age Influence Maternal Labor Supply and Child Care Response to the Earned Income Tax Credit?" *Journal of Labor Economics* 39(4): 895–929. doi:10.1086/711383.

National Institute for Early Education Research. 2025. "Vermont." *The State of Preschool 2024: State Preschool Yearbook*. New Brunswick, NJ: Rutgers University.

Olden, Andreas, and Jarle Møen. 2022. "The Triple Difference Estimator." *The Econometrics Journal* 25(3): 531–553.

Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe. 2023. "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature." *Journal of Econometrics* 235(2): 2218–2244.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics* 225(2): 175–199.

U.S. Census Bureau. 2010. *2010 American Community Survey Subject Definitions*. Washington, DC: U.S. Census Bureau. `https://www2.census.gov/programs-surveys/acs/tech_docs/subject_definitions/2010_ACSSubjectDefinitions.pdf`

U.S. Department of Health and Human Services, Administration for Children and Families, Office of Head Start. 2026. "How to Apply for Head Start Services." *HeadStart.gov*.

City of New York. 2014. *Ready to Launch: New York City's Implementation Plan for Free, High-Quality, Full-Day Universal Pre-Kindergarten*. New York: Office of the Mayor, Office of Management and Budget, Department of Education, and Administration for Children's Services.

Vermont General Assembly. 2014. "An Act Relating to Providing Access to Publicly Funded Prekindergarten Education." Act 166, H.270, signed May 28, 2014.

Figures

Figure 1: State-funded enrollment of four-year-olds, NIEER

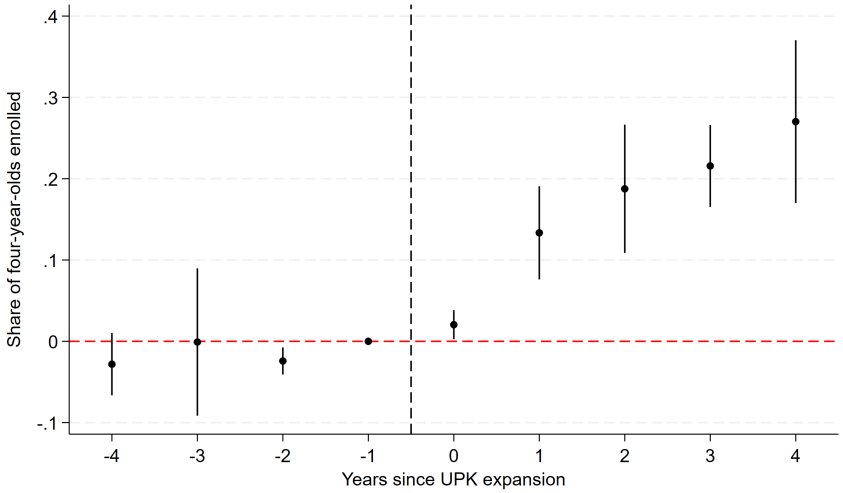
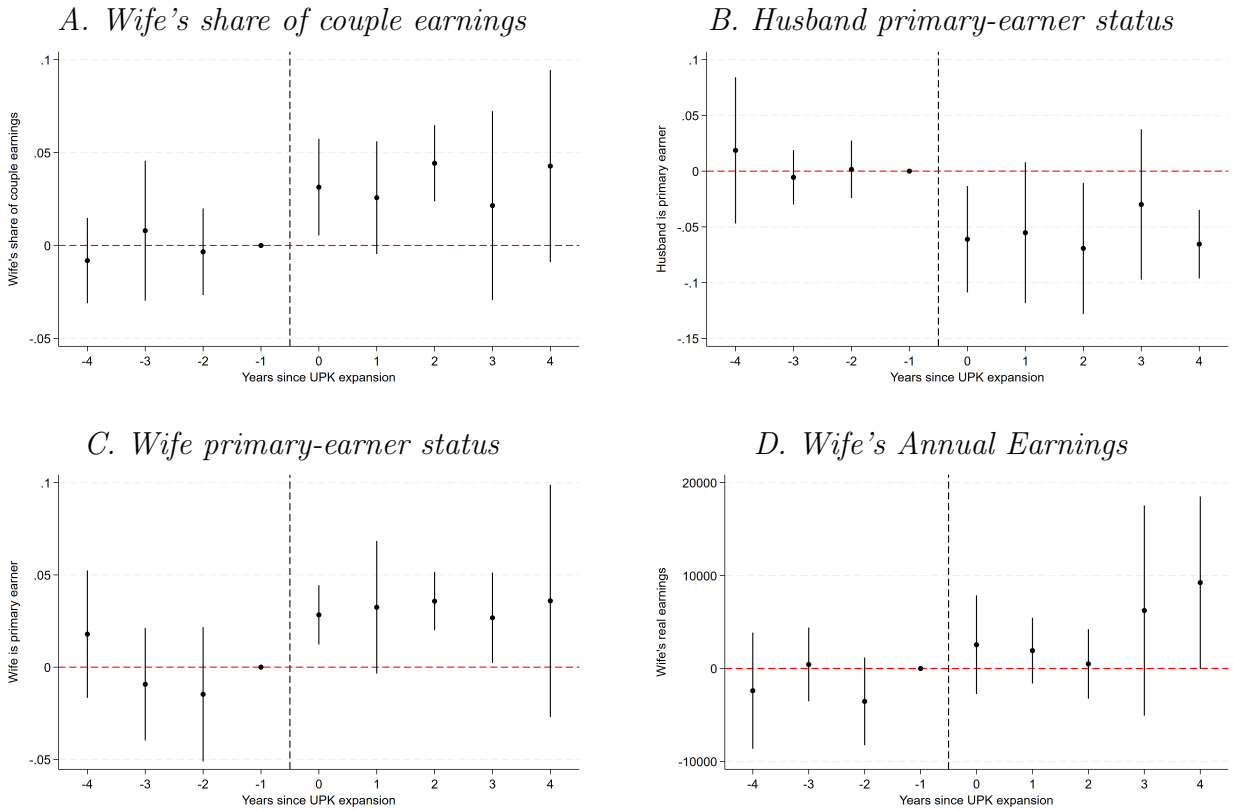
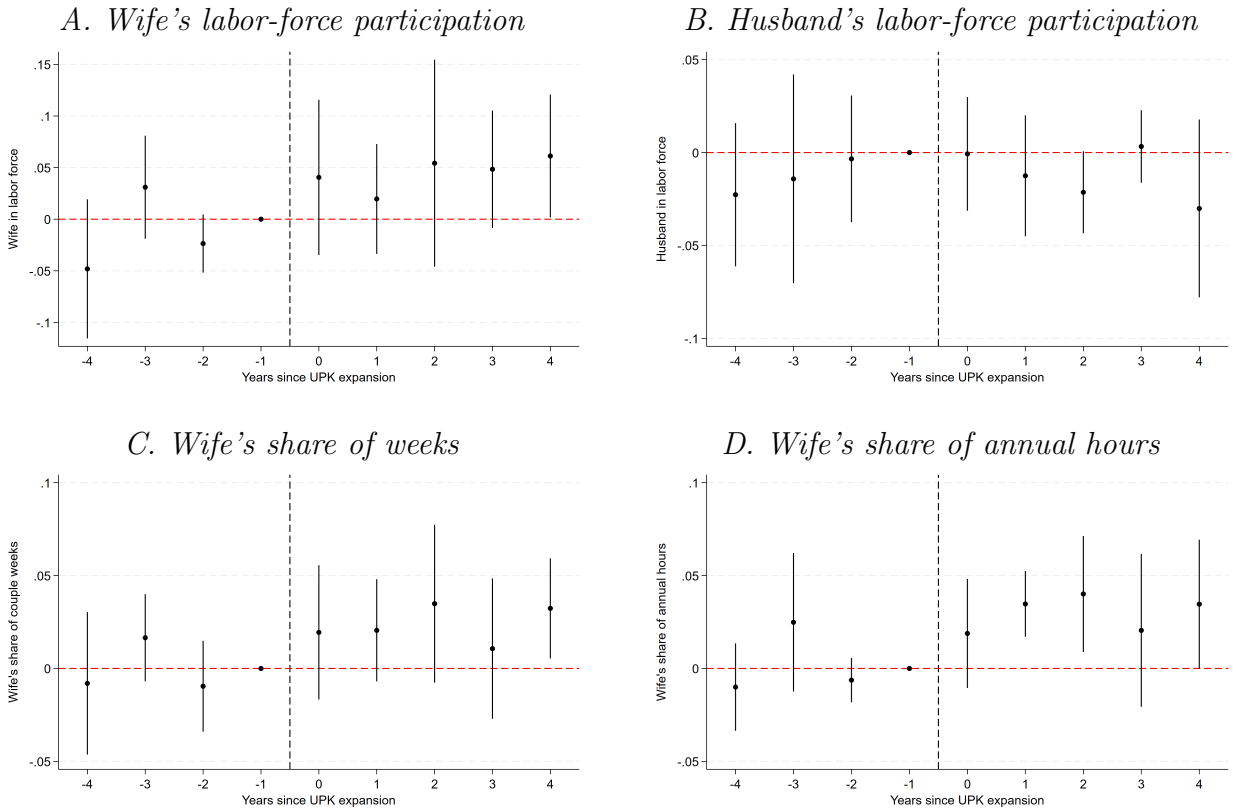


Figure 2: Event-study estimates: Division of Earnings Within Couples



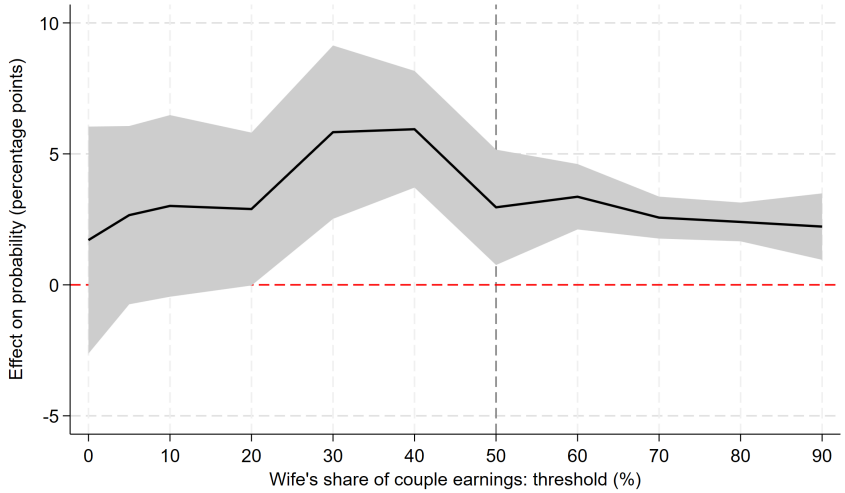
Notes: The figure reports event-study estimates from the preferred stacked triple-difference specification. Event year -1 is omitted. Vertical bars show 95 percent confidence intervals based on standard errors clustered at the jurisdiction level. Regressions are weighted using ACS household weights.

Figure 3: Event-study estimates: Labor supply and the Allocation of Market Work



Notes: The figure reports event-study estimates from the preferred stacked triple-difference specification. Event year −1 is omitted. Vertical bars show 95 percent confidence intervals based on standard errors clustered at the jurisdiction level. Regressions are weighted using ACS household weights.

Figure 4: Effects of UPK Across Wives' Earnings-Share Thresholds



Notes. The figure plots estimated effects of UPK on the probability that the wife's share of couple earnings exceeds each threshold, using the preferred stacked triple-difference specification. The shaded band represents 95% confidence intervals based on standard errors clustered by jurisdiction. The vertical dashed line marks equal earnings between spouses.

Tables

Table 1: Universal Pre-K Expansions

jurisdiction	Expansion year	Eligible ages	Pre-expansion enrollment	2019 enrollment	Main sample
Georgia	1995	4	–	60%	No
Oklahoma	1998	4	–	76%	No
West Virginia	2002	4	–	59%	No
Florida	2005	4	0%	75%	No
Iowa	2007	4	4%	66%	Yes
Wisconsin	2008	4	36%	72%	Yes
District of Columbia	2009	3–4	68%	87%	Yes
Vermont	2014	3–4	53%	78%	Yes
New York City	2014	4	–	–	Yes

Notes: UPK jurisdictions and treatment years follow Jackson, Turner, and Bastian (2025). Enrollment is the share of four-year-olds enrolled in state-funded pre-kindergarten. Pre-expansion enrollment is measured in the year before the listed UPK year where comparable data are available. Enrollment data begin in 2002, so pre-expansion rates are not shown for Georgia, Oklahoma, or West Virginia. The available state enrollment series does not separately report New York City, so NYC enrollment is not shown.

Table 2: Age-eligible households before expansion

	Treated jurisdictions	Never-treated jurisdictions
Labor supply		
Wife in labor force	0.649	0.637
Husband in labor force	0.946	0.949
Hours and weeks		
Wife, usual weekly hours	24.3	23.5
Husband, usual weekly hours	42.8	43.2
Wife, weeks worked	31.0	29.6
Husband, weeks worked	46.3	46.3
Earnings, 2010 dollars		
Wife	24,047	22,999
Husband	55,515	58,666
Couple	79,562	81,665
Division within the couple		
Wife's share of couple earnings	0.299	0.287
Wife's share of couple annual hours	0.306	0.296
Husband is primary earner	0.650	0.656
Wife is primary earner	0.147	0.139
Households	2,785	200,799

Notes. Means over the four years before expansion in each stack, weighted by ACS household weights. Event time is defined relative to each treated jurisdiction's expansion year, so the columns cover the same calendar years within each stack.

The treated column reproduces the pre-expansion means reported alongside the estimates in Tables 4 and 5. Households are counted once per survey year. Households in never-treated jurisdictions may appear in multiple stacks, so stacked observations exceed distinct households.

Table 3: First stage: enrollment and preschool attendance

	NIEER	ACS, estimation sample		
	Enrolled	Any	Public	Private
	(1)	(2)	(3)	(4)
Post × Treated	0.1788***	0.0262	0.0311***	−0.0049
	(0.0245)	(0.0176)	(0.0109)	(0.0086)
Pre-expansion mean	0.382	0.680	0.399	0.280
Observations	1,476	480,915	480,915	480,915
Unique household-years	—	167,892	167,892	167,892

Notes. Column 1 uses NIEER state-year data on enrollment of four-year-olds in state-funded pre-K. It includes Iowa, Wisconsin, the District of Columbia, and Vermont. New York State is excluded because enrollment is not reported separately for New York City.

Columns 2–4 use ACS data for households whose youngest child is four and include all five treated jurisdictions. All columns use a stacked double-difference specification. The ACS estimates are weighted by household weights. Standard errors are clustered by state in column 1 and by jurisdiction in columns 2–4.

The ACS observation count includes repeated observations across stacks while unique household-years counts each household-year record once, regardless of how many stacks it enters. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: The Division of Earnings Within Couples

	Wife's share of earnings	Husband is primary earner	Wife is primary earner	Wife's annual earnings
	(1)	(2)	(3)	(4)
ITT	0.0337***	−0.0594***	0.0336***	5,487**
	(0.0057)	(0.0114)	(0.0062)	(2,448)
RI p-value	0.008	0.022	0.012	0.140
Pre-expansion mean	0.299	0.650	0.147	24,047
Effect (% of mean)	11.3	-9.1	22.9	22.8
Observations	1,875,972	1,875,972	1,875,972	1,952,873
Unique household-years	746,663	746,663	746,663	776,667

Notes. Estimates from the stacked triple-difference specification, weighted by ACS household weights. All specifications include jurisdiction-by-year, child-age-by-year, and jurisdiction-by-child-age fixed effects, interacted with stack. Standard errors in parentheses are clustered by jurisdiction. Significance stars are based on clustered standard errors: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

RI p-values are from studentized randomization inference using 500 placebo assignments. Primary earner means receiving at least 60 percent of couple earnings. Earnings are measured in 2010 dollars.

Pre-expansion means refer to age-eligible households in treated jurisdictions before expansion. Effect (% of mean) is the estimated coefficient divided by the pre-expansion mean, multiplied by 100.

Observations are the number of records used in each regression, including repeated appearances across stacks. Unique household-years counts each household once, even if it appears in multiple stacks. Columns 1–3 are restricted to couples with positive total earnings.

Table 5: Labor Supply and the Allocation of Market Work

	Wife in labor force (1)	Husband in labor force (2)	Wife's share of weeks (3)	Wife's share of annual hours (4)
ITT	0.0549***	−0.0020	0.0236***	0.0273***
	(0.0122)	(0.0057)	(0.0065)	(0.0051)
RI p-value	0.028	0.776	0.038	0.010
Pre-expansion mean	0.649	0.946	0.348	0.306
Effect (% of mean)	8.5	-0.2	6.8	8.9
Observations	1,952,873	1,952,873	1,930,230	1,930,230
Unique household-years	776,667	776,667	768,086	768,086

Notes. Estimates from the stacked triple-difference specification. Weighting, fixed effects, and clustering are as described in Table 4. Standard errors in parentheses are clustered by jurisdiction. Significance stars are based on clustered standard errors: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

RI p-values are from studentized randomization inference using 500 placebo assignments. Annual hours are estimated using usual weekly hours multiplied by weeks worked.

Pre-expansion means refer to age-eligible households in treated jurisdictions before expansion. Effect (% of mean) is the estimated coefficient divided by the pre-expansion mean, multiplied by 100.

Columns 3 and 4 are restricted to couples with positive total weeks or annual hours, respectively.

Table 6: Effect on the Probability That the Wife's Earnings Share Exceeds c

c	Estimate	(s.e.)	Pre-expansion mean
0.00	0.0171	(0.0221)	0.678
0.05	0.0266	(0.0174)	0.639
0.10	0.0301*	(0.0177)	0.610
0.20	0.0289*	(0.0149)	0.554
0.30	0.0583***	(0.0169)	0.455
0.40	0.0594***	(0.0114)	0.350
0.50	0.0296**	(0.0112)	0.227
0.60	0.0336***	(0.0064)	0.143
0.70	0.0257***	(0.0041)	0.104
0.80	0.0240***	(0.0038)	0.086
0.90	0.0222***	(0.0065)	0.076

Notes. Each row reports the estimated effect of UPK on the probability that the wife's share of couple earnings exceeds c, using the preferred stacked triple-difference specification. Estimates are weighted by ACS household weights, with standard errors clustered by jurisdiction.

The $c = 0.40$ estimate is the negative of the husband-primary-earner estimate in Table 4. The $c = 0.60$ estimate also matches the wife-primary-earner estimate to four decimal places, although the latter includes couples earning exactly 60% of couple earnings.

Pre-expansion means refer to eligible households in treated jurisdictions before expansion. Figure A3 reports event studies for the thresholds from 0.20 to 0.70, which span the range over which the estimates change. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Appendix Tables

Table A1: All four groups before expansion

	Age-eligible		Youngest is 6–8		Difference
	Treated	Never-treated	Treated	Never-treated	in gaps
Labor supply					
Wife in labor force	0.649	0.637	0.753	0.703	−0.038
Husband in labor force	0.946	0.949	0.944	0.944	−0.003
Both spouses earning	0.587	0.569	0.633	0.623	0.008
Hours and weeks					
Wife, usual weekly hours	24.27	23.52	27.38	25.80	−0.83
Husband, usual weekly hours	42.83	43.19	42.77	43.11	−0.03
Wife, weeks worked	30.97	29.62	35.00	32.54	−1.12
Husband, weeks worked	46.30	46.35	45.90	46.18	0.23
Earnings, 2010 dollars					
Wife	24,047	22,999	27,128	24,795	−1,285
Husband	55,515	58,666	56,236	62,007	2,619
Couple	79,562	81,665	83,364	86,801	1,334
Division within the couple					
Wife's share of couple earnings	0.299	0.287	0.341	0.309	−0.020
Wife's share of couple annual hours	0.306	0.296	0.351	0.324	−0.017
Husband is primary earner	0.650	0.656	0.599	0.641	0.036
Wife is primary earner	0.147	0.139	0.175	0.153	−0.014
Characteristics					
Wife's age	35.45	34.84	38.88	38.52	0.24
Husband's age	38.14	37.27	41.39	40.87	0.35
Wife has a four-year degree	0.386	0.376	0.346	0.359	0.023
Husband has a four-year degree	0.346	0.354	0.319	0.351	0.024
Wife is white	0.722	0.765	0.743	0.775	−0.011
Number of children	2.299	2.266	2.297	2.274	0.010
Stacked observations	2,785	308,609	6,562	538,613	
Households	2,785	200,799	6,562	260,311	

Notes. Means over the four years before expansion in each stack, weighted by ACS household weights. The four columns cover the same calendar years within each stack. Age-eligible households have a youngest child aged four, or three or four in the District of Columbia and Vermont. The comparison group has a youngest child aged six to eight.

The final column is the age-eligible minus older-child difference in treated jurisdictions, less the corresponding difference in never-treated jurisdictions, computed from unrounded means. The triple-difference design does not require this pre-expansion difference to be zero. Identification instead requires parallel trends in the age-group differences in the absence of UPK expansion.

A household in a never-treated jurisdiction enters every stack whose window contains its survey year, so stacked observations exceed distinct households in those columns and coincide with them in the treated columns.

Table A2: Effects of UPK on Women's Labor Supply by Marital Status

	All women	Married	Married, in couples	Cohabiting	Single
	(1)	(2)	(3)	(4)	(5)
Labor-force participation	0.0391***	0.0580***	0.0562***	0.0338	0.0085
	(0.0131)	(0.0143)	(0.0133)	(0.0224)	(0.0117)
Employment	0.0306***	0.0533***	0.0524***	0.0380	−0.0131
	(0.0096)	(0.0080)	(0.0085)	(0.0353)	(0.0139)
Pre-expansion means					
Labor-force participation	0.694	0.647	0.652	0.746	0.776
Employment	0.632	0.608	0.614	0.670	0.669
Observations	2,844,976	2,011,876	1,958,285	193,664	639,435
Distinct person-years	1,133,216	801,229	778,784	79,013	252,974

Notes.

Estimates use the stacked triple-difference specification on ACS person-level data for women aged 18–55 whose youngest own child is age-eligible or in the comparison group. All regressions use person weights, with standard errors in parentheses clustered by jurisdiction.

Pre-expansion means refer to age-eligible women in treated jurisdictions before expansion. Observations include repeated appearances across stacks; unique women are counted once per survey year. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A3: Alternative Specifications

	Stacked		TWFE	
	Triple diff. (1)	Double diff. (2)	Triple diff. (3)	Double diff. (4)
Wife in labor force	0.0549	0.0511	0.0537	0.0481
	(0.0122)	(0.0146)	(0.0122)	(0.0152)
Wife's earnings share	0.0337	0.0271	0.0326	0.0269
	(0.0057)	(0.0052)	(0.0057)	(0.0052)
Husband is primary earner	-0.0594	-0.0525	-0.0566	-0.0513
	(0.0114)	(0.0062)	(0.0107)	(0.0060)
Wife is primary earner	0.0336	0.0228	0.0379	0.0290
	(0.0062)	(0.0081)	(0.0031)	(0.0056)
Wife's share of weeks	0.0236	0.0232	0.0214	0.0205
	(0.0065)	(0.0081)	(0.0064)	(0.0083)
Wife's share of annual hours	0.0273	0.0254	0.0257	0.0237
	(0.0051)	(0.0057)	(0.0052)	(0.0061)
Wife's real earnings	5,487	5,579	5,279	5,390
	(2,448)	(1,872)	(2,450)	(1,943)
Child-age comparison	Yes	No	Yes	No
Jurisdiction-year effects	Yes	No	Yes	No

Notes. Column 1 is the preferred specification. Standard errors in parentheses are clustered by jurisdiction. TWFE specifications retain observations outside the four-year event window using endpoint indicators.

Table A4: Sensitivity to Excluding Each Treated Jurisdiction

	Earnings and Primary Earner Status				Labor Supply and Market Work				Share of treated (%)
Jurisdiction excluded	Wife's earnings share	Husband primary earner	Wife primary earner	Wife's annual earnings	Wife in labor force	Husband in labor force	Wife's share of weeks	Wife's share of hours	
None	0.0337	−0.0594	0.0336	5,487	0.0549	−0.0020	0.0236	0.0273	—
	(0.0057)	(0.0114)	(0.0062)	(2,448)	(0.0122)	(0.0057)	(0.0065)	(0.0051)	
New York City	0.0307	−0.0538	0.0325	2,454	0.0582	−0.0074	0.0209	0.0255	33.4
	(0.0078)	(0.0152)	(0.0103)	(1,381)	(0.0195)	(0.0059)	(0.0099)	(0.0079)	
Wisconsin	0.0355	−0.0684	0.0299	7,111	0.0438	0.0005	0.0234	0.0265	35.0
	(0.0082)	(0.0149)	(0.0088)	(2,921)	(0.0121)	(0.0079)	(0.0098)	(0.0075)	
Iowa	0.0388	−0.0664	0.0319	6,858	0.0661	0.0027	0.0299	0.0322	19.9
	(0.0049)	(0.0132)	(0.0076)	(2,576)	(0.0095)	(0.0050)	(0.0050)	(0.0037)	
Vermont	0.0312	−0.0530	0.0387	4,980	0.0524	−0.0022	0.0199	0.0246	7.3
	(0.0053)	(0.0094)	(0.0029)	(2,564)	(0.0126)	(0.0060)	(0.0059)	(0.0047)	
District of Columbia	0.0328	−0.0574	0.0332	5,662	0.0537	−0.0038	0.0241	0.0278	4.3
	(0.0057)	(0.0112)	(0.0064)	(2,494)	(0.0125)	(0.0057)	(0.0067)	(0.0052)	

Notes. Each row reports estimates from the preferred stacked triple-difference specification, excluding the indicated treated jurisdiction and its corresponding stack. The first row reports the full-sample estimates. All regressions are weighted using ACS household weights. Standard errors in parentheses are clustered by jurisdiction.

Share of treated is the excluded jurisdiction's percentage of age-eligible households in treated jurisdictions after expansion.

Table A5: Selection into the Estimation Sample

	Age-Restricted Sample			Broader Child-Age Sample	
	Married, spouse present	Cohabiting	Number of children	Youngest child aged four	Number of children
	(1)	(2)	(3)	(4)	(5)
ITT	0.0105	−0.0079	−0.0260*	−0.0014	−0.0060
	(0.0111)	(0.0086)	(0.0148)	(0.0015)	(0.0105)
Pre-expansion mean	0.611	0.085	2.156	0.102	2.101
Specification	DDD	DDD	DDD	DiD	DiD

Notes. The sample includes women aged 18–55 regardless of marital status. Columns 1–3 use the stacked triple-difference specification, comparing women whose youngest child is UPK-age-eligible with women whose youngest child is aged six to eight. Columns 4–5 use a stacked double-difference specification including women whose youngest child is aged zero to eight. These columns examine changes in the age distribution of children and family size that could affect entry into the age-eligible sample.

All regressions use ACS person weights. Standard errors in parentheses are clustered by jurisdiction. Pre-expansion means refer to women in treated jurisdictions before expansion. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A6: Could Changes in Sample Composition Explain the Main Results?

	Estimated	Existing	Required entrant mean		Percentile
Outcome	effect	mean	$s = 1.7\%$	$s = 5.0\%$	
Division of earnings					
Wife's share of couple earnings	0.0337	0.299	2.294	0.971	92.7
Husband is primary earner	-0.0594	0.650	-2.867	-0.535	—
Wife is primary earner	0.0336	0.147	2.136	0.817	—
Labor supply					
Wife in labor force	0.0549	0.649	3.900	1.744	—
Wife's share of couple weeks	0.0236	0.348	1.745	0.819	95.6
Wife's share of couple annual hours	0.0273	0.306	1.922	0.851	95.7

Notes. The table reports the mean outcome among entering couples required for changes in sample composition alone to explain each estimated effect. The required mean is $m = \mu + \beta/s$, where μ is the mean among existing couples, β is the estimated effect, and s is the share of entrants in the post-expansion married sample.

The estimated effect on marriage is 0.0105 (s.e. 0.0111), against a pre-expansion mean of 0.611 (Table A5). The implied entrant share is $0.0105/(0.611 + 0.0105) = 1.7\%$. Using the upper 95% confidence limit gives 5.0%. Both calculations assume that the change in marriage reflects entry without offsetting exits.

All outcomes are bounded between zero and one. The final column reports the percentile of the required entrant mean in the observed distribution of each continuous outcome under the larger entrant-share scenario ($s = 5.0\%$).

Table A7: Sample Composition and Sensitivity to Demographic Controls

Panel A. Changes in household characteristics

	Wife's age (1)	Husband's age (2)	Wife has college degree (3)	Husband has college degree (4)	Wife is white (5)
ITT	0.334***	0.334***	0.025*	0.011	0.004
	(0.085)	(0.061)	(0.012)	(0.023)	(0.015)
Pre-expansion mean	35.45	38.14	0.386	0.346	0.722

Panel B. Sensitivity to demographic controls

	Baseline (1)	+ Age (2)	+ Education, race (3)	Change (%) (4)
Wife's earnings share	0.0337***	0.0336***	0.0319***	−5.3
	(0.0057)	(0.0058)	(0.0083)	
Husband is primary earner	−0.0594***	−0.0593***	−0.0564***	−5.0
	(0.0114)	(0.0115)	(0.0142)	
Wife is primary earner	0.0336***	0.0324***	0.0311***	−7.3
	(0.0062)	(0.0061)	(0.0067)	
Wife in labor force	0.0549***	0.0563***	0.0540***	−1.7
	(0.0122)	(0.0122)	(0.0141)	
Wife's share of weeks	0.0236***	0.0239***	0.0226***	−4.3
	(0.0065)	(0.0067)	(0.0080)	
Wife's share of annual hours	0.0273***	0.0276***	0.0263***	−3.8
	(0.0051)	(0.0053)	(0.0069)	
Wife's real earnings	5,487**	5,279**	4,832**	−11.9
	(2,448)	(2,431)	(2,392)	

Notes. Panel A examines changes in household characteristics following UPK expansion and uses 1,952,873 stacked observations representing 776,667 unique households. Panel B examines the sensitivity of the main estimates to demographic controls. Column 2 adds age and age squared for both spouses; column 3 also adds education and race. Change is the percentage change relative to the baseline estimate. The estimation sample is unchanged when demographic controls are added. Sample sizes for each outcome are reported in Tables 4 and 5.

All regressions use the preferred stacked triple-difference specification and ACS household weights. Standard errors in parentheses are clustered by jurisdiction. Pre-expansion means refer to eligible households in treated jurisdictions before expansion. Sample selection is examined separately in Table A5. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A8: Restricting the Comparison Group by State-Funded Pre-K Enrollment

	All	Comparison enrollment at or below			
	controls	25%	10%	5%	0%
	(1)	(2)	(3)	(4)	(5)
Panel A. Estimated effects					
Wife's share of couple earnings	0.0337	0.0341	0.0402	0.0410	0.0401
	(0.0057)	(0.0065)	(0.0055)	(0.0058)	(0.0081)
Husband is primary earner	−0.0594	−0.0610	−0.0616	−0.0585	−0.0632
	(0.0114)	(0.0135)	(0.0118)	(0.0127)	(0.0196)
Wife is primary earner	0.0336	0.0331	0.0447	0.0428	0.0389
	(0.0062)	(0.0064)	(0.0103)	(0.0093)	(0.0097)
Wife in labor force	0.0549	0.0537	0.0614	0.0589	0.0593
	(0.0122)	(0.0131)	(0.0141)	(0.0133)	(0.0156)
Husband in labor force	−0.0020	−0.0021	−0.0052	−0.0078	−0.0100
	(0.0057)	(0.0063)	(0.0078)	(0.0086)	(0.0100)
Wife's share of couple weeks	0.0236	0.0238	0.0272	0.0257	0.0236
	(0.0065)	(0.0071)	(0.0063)	(0.0066)	(0.0083)
Wife's share of annual hours	0.0273	0.0276	0.0320	0.0329	0.0316
	(0.0051)	(0.0058)	(0.0054)	(0.0054)	(0.0071)
Wife's real earnings	5,487	5,680	5,842	6,105	5,363
	(2,448)	(2,475)	(1,709)	(1,712)	(1,741)
Panel B. Largest absolute pre-expansion t-statistic					
Wife's share of couple earnings	0.71	0.48	0.71	0.28	2.50
Husband is primary earner	0.57	0.48	0.44	0.40	1.60
Wife is primary earner	1.04	1.23	1.70	1.94	1.91
Wife in labor force	1.68	1.37	1.65	1.88	2.13
Husband in labor force	1.19	1.42	0.96	0.99	0.42
Wife's share of couple weeks	1.42	1.65	1.36	0.94	1.79
Wife's share of annual hours	1.34	1.52	1.41	1.06	1.62
Wife's real earnings	1.51	1.55	2.02	1.90	1.89
Comparison jurisdictions	41	30	19	15	9
Jurisdictions	46	35	24	20	14

Notes. Estimates from the preferred stacked triple-difference specification, weighted by ACS household weights. Standard errors in parentheses are clustered by jurisdiction. Column 1 is the baseline specification.

Comparison jurisdictions are restricted according to their maximum state-funded pre-K enrollment rate during each stack's nine-year window, using NIEER data. Restrictions are applied separately within each stack. The six jurisdictions excluded from the baseline specification remain excluded. New York State outside New York City is excluded from the restricted comparison groups because its enrollment rate cannot be separated from New York City.

Panel B reports the largest absolute t-statistic among the three pre-expansion event-study coefficients. Jurisdiction counts refer to jurisdictions appearing in at least one stack.

Appendix Figures

Figure A1: Public Preschool Attendance of four-year-olds, ACS sample

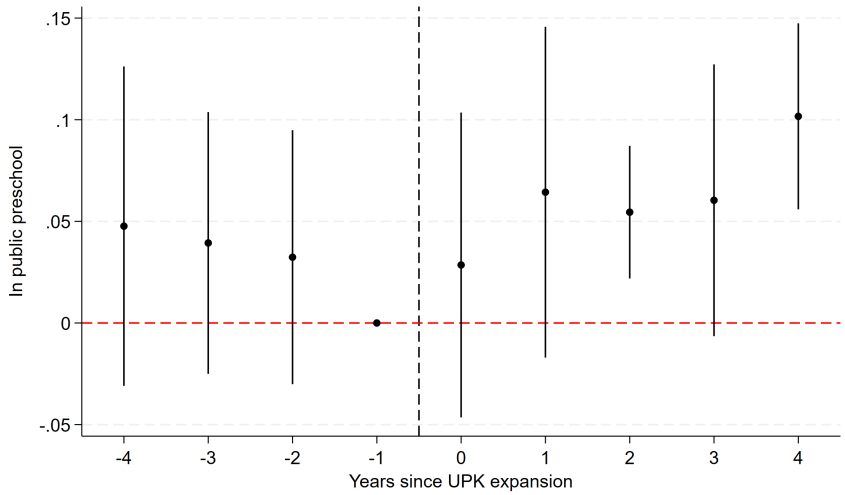
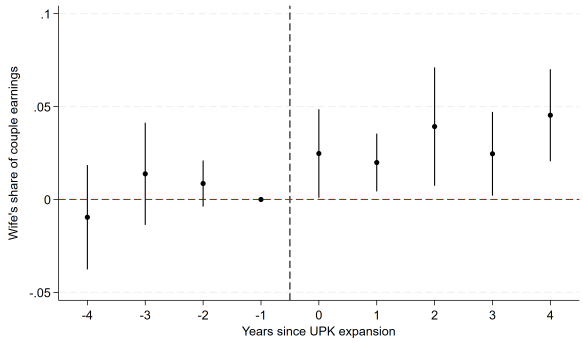
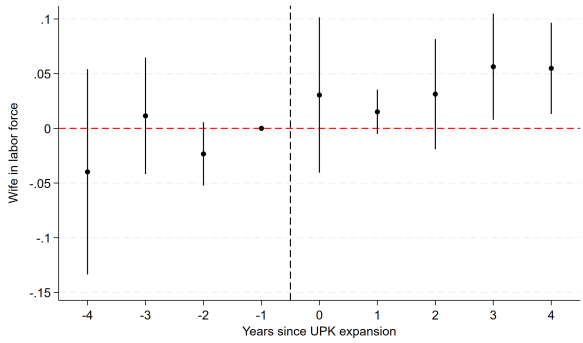


Figure A2: Stacked Double-Difference Event-Study Estimates

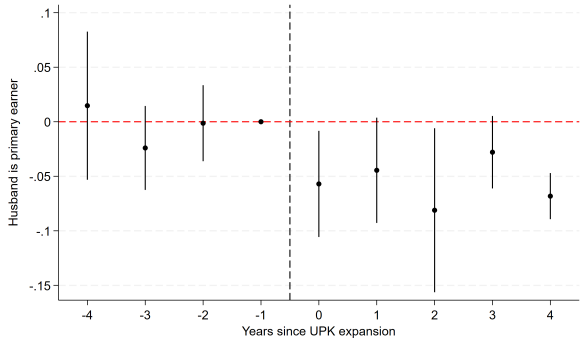
(a) Wife's earnings share



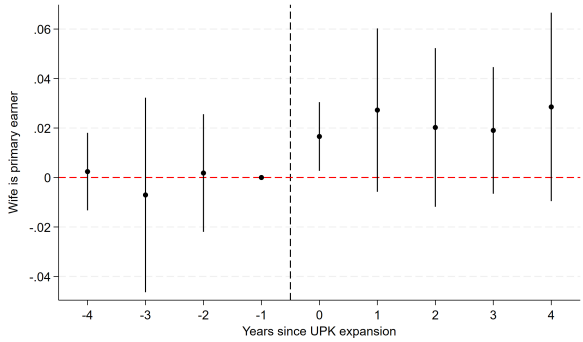
(b) Wife in labor force



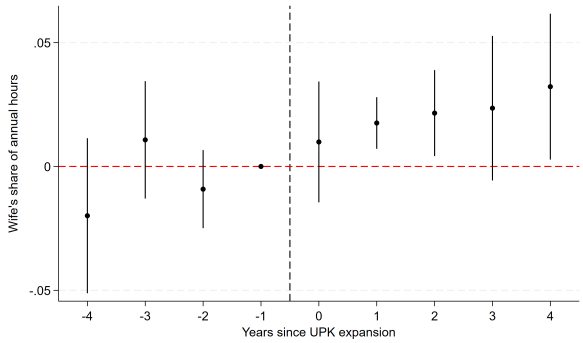
(c) Husband is primary earner



(d) Wife is primary earner



(e) Wife's share of annual hours



(f) Husband in labor force

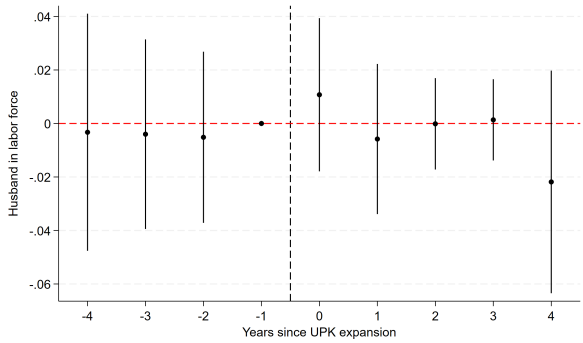
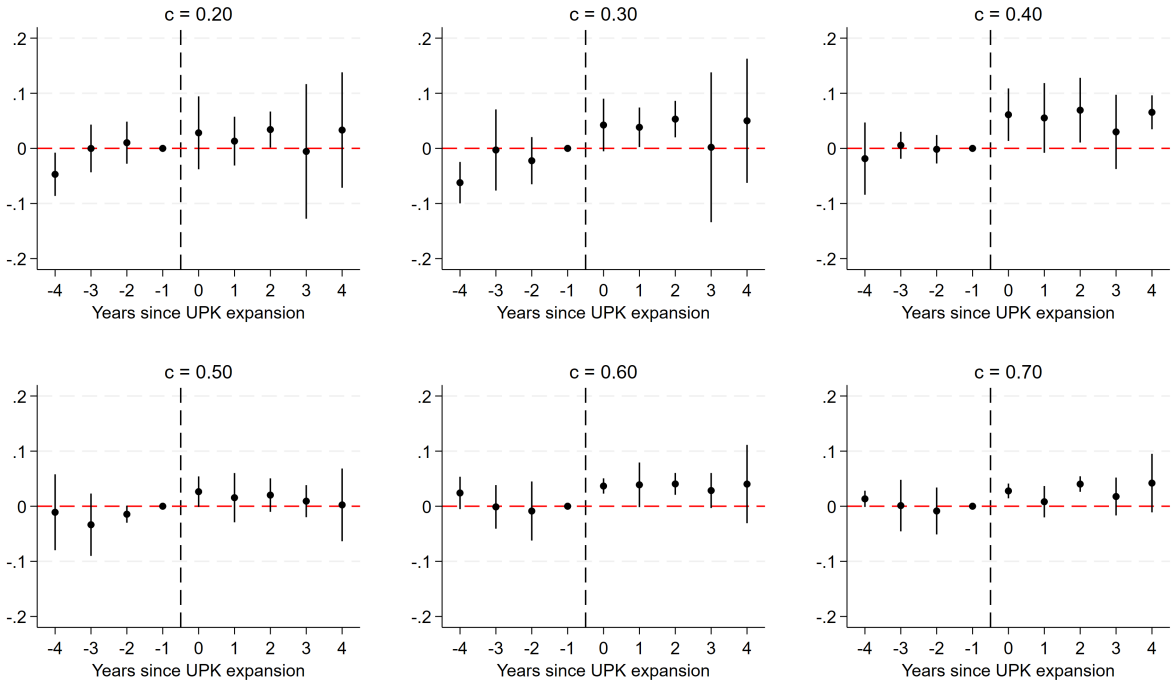


Figure A3: Event-Study Estimates Across Wives' Earnings-Share Thresholds



Notes. Each panel reports event-study estimates of the effect of UPK on the probability that the wife's share of couple earnings exceeds the indicated threshold, using the preferred stacked triple-difference specification. Event year -1 is the omitted reference period. Vertical bars represent 95% confidence intervals based on standard errors clustered by jurisdiction.