ROCKWOOL Foundation Berlin
Institute for the Economy and the Future of Work (RFBerlin)

DISCUSSION PAPER SERIES 162/26

Effect of Remote Work on the Child Penalty: Evidence from the United States

Ahmet Gulek, Christina Langer

www.rfberlin.com

JUNE 2026

Effect of Remote Work on the Child Penalty: Evidence from the United States

Authors

Ahmet Gulek, Christina Langer

Reference

JEL Codes: J13, J16, J22, J31

Keywords: Working from Home, Child Penalty, Gender Inequality

Recommended Citation: Ahmet Gulek, Christina Langer (2026): Effect of Remote Work on the Child Penalty: Evidence from the United States. RFBerlin Discussion Paper No. 162/26

Imprint

RFBerlin
ROCKWOOL Foundation Berlin –
Institute for the Economy
and the Future of Work

Gormannstrasse 22, 10119 Berlin
Tel: +49 (0) 151 143 444 67
E-mail: info@rfberlin.com
Web: www.rfberlin.com

Effect of Remote Work on the Child Penalty: Evidence from the United States

Ahmet Gulek* Christina Langer

June 6, 2026

Most recent draft here.

Abstract

We study the effect of the post-COVID expansion of remote work on the child penalty. Our empirical design proceeds in three steps: using pseudo-panels to estimate child penalties across occupations, exploiting cross-occupational variation in remote work adoption, and using synthetic controls to adjust for differential pre-COVID trajectories. Remote work reduces the hours penalty for mothers: each percentage point increase in an occupation's remote work share raises mothers' hours worked by 0.23%, eliminating one-sixth of the baseline motherhood penalty for the average remote occupation. We find no effects on employment, income, or wages for mothers, and no average effects on men.

Keywords: Working from Home, Child Penalty, Gender Inequality

JEL codes: J13, J16, J22, J31

*Gulek: Postdoctoral Scholar, Immigration Policy Lab, Stanford University; RFBerlin (e-mail: agulek@stanford.edu). Langer: Ifo Institute; Digital Fellow, Stanford Digital Economy Lab (e-mail: langerc@stanford.edu). We are thankful to Emma Harrington, Joanna Venator and the participants at the SOLE 2026 Annual Meeting for their invaluable comments.

1 Introduction

Since the 1970s, the gender gap in earnings and employment in the United States has narrowed substantially (Goldin, 2024). Nevertheless, significant disparities persist: women participate less in the labor force than men and, when employed, earn lower wages. A growing literature shows that these gendered disparities are largely explained by the child penalty: the differential impact of parenthood on labor market outcomes between women and men (Kleven et al., 2019, 2024a; Cortés and Pan, 2023). Further progress in gender equality likely requires changes in child penalties, yet traditional family policies, including parental leave, childcare subsidies, and flexible scheduling, have had limited effects, even in Scandinavia (Kleven et al., 2024b; Olivetti and Petrongolo, 2017). We study whether the expansion of remote work after the COVID-19 pandemic has affected the child penalty in the United States. Roughly 30 percent of US workdays now occur remotely (Aksoy et al., 2025), making this one of the largest natural experiments in labor markets in recent decades.

Theory alone cannot sign the impact of remote work on child penalties. Remote work provides flexibility and reduces commuting time, helping mothers who bear a disproportionate share of caregiving balance work and family (Goldin, 2014; Goldin and Katz, 2016). Mothers who would otherwise reduce their hours after childbirth may find it easier to maintain their pre-birth schedule, which would lower the child penalty. Conversely, if mothers disproportionately select into remote work for its non-wage amenity value (Mas and Pallais, 2017), the result could be lower earnings and reduced career progression (Bloom et al., 2015). A third possibility is that working from home while caring for young children lowers hourly productivity, so that even if mothers work more hours, their output and earnings do not increase proportionally. Experimental evidence supports this channel: women randomly assigned to work from home complete tasks more slowly as domestic responsibilities fragment work sessions (Ho et al., 2025). These forces oppose each other, making the net impact of remote work on the child penalty an empirical question.

Three challenges arise. First, estimating child penalties requires panel data, which are limited in size. The PSID and NLSY track only a few thousand parents; estimating child penalties separately for each of 94 occupations requires orders of magnitude more observations. Second, while public datasets like the ACS can measure current remote work arrangements, precise estimates of the treatment require tracking remote work incidence *before* COVID across occupations, information that standard surveys do not provide. Third, the adoption of remote work is not random across occupations. Remote-friendly occupations, disproportionately high-skill, white-collar jobs, were already on differential child penalty trajectories before the pandemic, making the standard parallel trends assumption unlikely to hold.[1] As we document, the employment child penalty in eventually-remote occupations was shrinking throughout 2001–2019.

We address these challenges in three steps. First, we use the pseudo-panel approach of Kleven (2025), which matches parents to observationally similar non-parents in the ACS (∼3 million observations per year), providing sufficient statistical power for occupation-level estimation. We exclude the 2020 and 2021 ACS to avoid bias from pandemic-specific disruptions unrelated to remote work. Second, we use Lightcast (formerly Burning Glass Technologies), which classifies the universe of US online job postings as remote or on-site since 2015, to measure the change in remote work share for each occupation before and after COVID. Third, we apply synthetic control difference-in-differences (SC-DiD) to construct counterfactual trends for each occupation, adjusting for the differential pre-COVID trajectories that would bias standard DiD.

Our main result is that remote work reduces the hours child penalty for mothers. The effect on income is positive and similar in magnitude to the hours effect but not statistically significant in the preferred specification, and not robust across specifications. Employment and wage effects are indistinguishable from zero across all specifications. Each additional percentage point of remote work raises mothers' log hours by 0.23 percent ($p = 0.029$), implying approximately 2 percent more hours for the average remote occupation. This effect eliminates roughly one-sixth of the baseline hours penalty for mothers in remote occupations, an economically meaningful reduction given that traditional family policies have had limited effects on child penalties (Kleven et al., 2024b). The hours effect operates broadly across the distribution, shifting mothers from part-time toward full-time schedules. We find no average effects on men. That income and wages do not rise commensurately points to a productivity offset: evidence from the American Time Use Survey shows that mothers in remote occupations sharply increased the time spent supervising children while working after COVID, consistent with lower hourly productivity absorbing the hours gains. These results are robust across specifications, placebo backtests, and CPS rotating-panel replication.

[1] Possible explanations include the diffusion of communication technologies that disproportionately benefitted knowledge-work occupations, making them incrementally more compatible with flexible work arrangements even before remote work expanded broadly.

We also find no evidence of anticipatory occupational sorting.

Our results speak to two literatures. First, a growing body of work documents that child penalties are notoriously unresponsive to government policies: expansions of parental leave, childcare subsidies, and gender quotas have had limited effects on the earnings gap between mothers and fathers (Kleven et al., 2024b; Olivetti and Petrongolo, 2017). A notable exception is Ciasullo and Uccioli (2023), who provide the first causal evidence that work arrangements affect child penalties: exploiting Australia's 2009 Fair Work Act, which entitled mothers to request reduced hours on predictable, permanent contracts, they find an 8 pp increase in labor force participation and a 22 percent reduction in the hours penalty. We provide new evidence that remote work, a market-driven change in work arrangements rather than a government mandate, also reduces the hours child penalty in the United States.

Second, we provide occupation-level estimates of how remote work affects child penalties, contributing to the literature on remote work and labor markets (Bloom et al., 2015, 2022; Emanuel and Harrington, 2024; Aksoy et al., 2025). Most closely related to our work, three concurrent papers study the relationship between remote work and child penalties. Harrington and Kahn (2026) find a positive correlation between pre-COVID remote work adoption and the narrowing of the mother–non-mother employment gap. Our approach complements theirs in three ways: we estimate causal child penalties from event studies rather than raw gaps, we adjust for differential pre-COVID trajectories via synthetic controls (after which the employment association disappears), and we document that the effect operates on the intensive margin. Scott and Sundberg (2025) find null effects in Sweden, while Zarate (2025) finds a 4 pp reduction in the employment penalty in Latin America. Relative to all three papers, we adjust for differential pre-COVID trajectories via synthetic controls, document a positive effect on hours, and identify a productivity mechanism that can explain why these additional hours do not translate into significant earnings gains.

2 Data

2.1 Remote work

We use Lightcast (formerly Burning Glass Technologies), which covers nearly all US online job postings (∼40 million per year) provided by Lightcast and Hansen et al. (2023). They classify each posting as remote or on-site based on text analysis of the job description. Unlike worker-reported surveys, this measures the *demand* side of remote work: the share of positions that employers advertise as remote. We aggregate to 94 three-digit SOC occupations and compute the remote work share in 2017–2019 (baseline) and 2023–2025 (post).[2] Our treatment variable is ΔR_o, the percentage-point change in remote work share for occupation o between these two periods. We also classify occupations as remote or non-remote based on whether ΔR_o is above or below the median.[3] Appendix Figure A.14 validates this measure against realized telework: across 94 occupations, ΔR_o is strongly correlated ($r = 0.82$) with the share of workers who report teleworking in the 2025 CPS.

Figure 1 documents the variation in remote work adoption that identifies our results. Panel (a) plots the monthly remote work share for remote and non-remote occupations from 2015 through 2025. Before COVID, both groups were essentially flat at low levels. After March 2020, remote occupations experienced a sharp, persistent increase while non-remote occupations barely moved. The pre-COVID stability confirms the pandemic drove adoption. Panel (b) shows the cross-sectional variation underlying this aggregate pattern. Remote work adoption is positively correlated with skill intensity: low-skill occupations saw minimal increases, while high-skill occupations like computer, legal, and marketing experienced increases exceeding 20 pp. The treatment variation has two important features: there is large heterogeneity across occupations in employer-offered remote work, and remoteness is not randomly assigned. Remote and non-remote occupations differ systematically in skill intensity and likely in other unobserved characteristics correlated with child penalty trends.

2.2 Labor market outcomes

ACS. Our primary dataset is the ACS, 2001–2024 (∼3 million observations per year). We restrict the sample to individuals aged 25–55 who had their first child between ages 25 and 45. Following Kleven (2025), we construct pseudo-panels by matching parents to observationally similar

[2]We exclude 2020 from all analyses to avoid contamination from the acute phase of the pandemic. In the ACS, the post period begins in 2022 (excluding 2020–2021). In the CPS, the post period begins in 2021. The difference arises because ACS employment questions refer to the prior year, while CPS questions refer to the past week. Note that ACS "usual hours worked per week" refers to a typical week at the time of the survey, while employment and income refer to the prior year. Excluding the 2021 ACS is therefore conservative for the hours variable, which captures 2021 conditions rather than 2020, but we do so to maintain a consistent sample across outcomes.

[3]Lightcast measures *posted* remote work availability, not realized remote work arrangements. If measurement error in ΔR_o is classical (uncorrelated with the true remote share and with child penalty changes), it attenuates the continuous estimates, implying that the true effect per percentage point of actual remote work could be larger than we report. If the gap between posted and realized remote work varies systematically across occupations, the bias could go in either direction.

Figure 1: Remote Work Adoption by Occupation

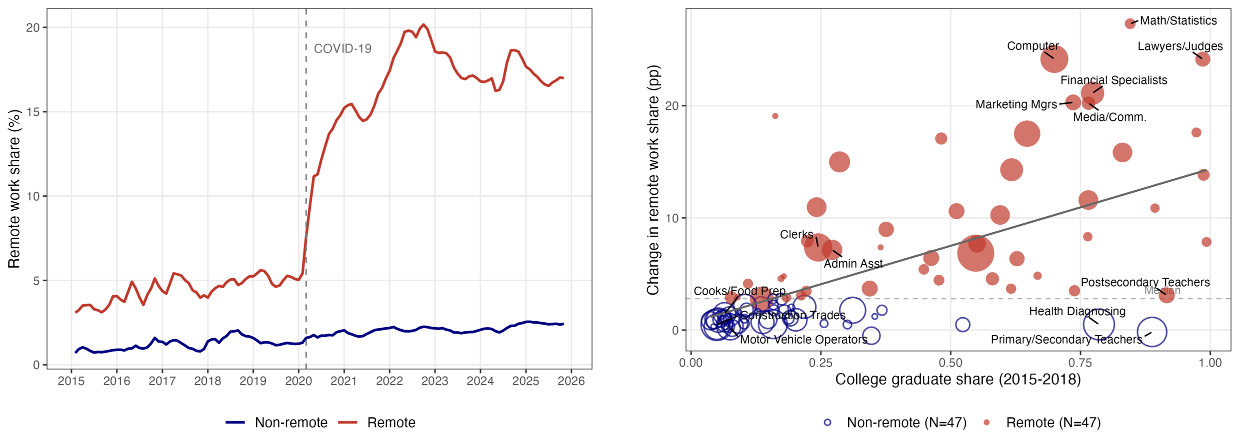

(a) Remote work share over time

(b) Occupation skill and remote work adoption

Panel (a): monthly share of job postings offering remote work, 3-month moving average. Red: remote occupations ($\Delta R_o \geq$ median); blue: non-remote. 2015–2018 from Lightcast extract, level-adjusted to match at 2019; 2019–2025 from Hansen et al. (2023). Dashed line: March 2020. Panel (b): each bubble is a 3-digit SOC occupation, sized by employment. x-axis: college graduate share (ACS, 2015–2018). y-axis: ΔR_o, change in remote work share, 2017–19 vs. 2023–25. Dashed line: median $\Delta R_o = 2.8$ pp. Red: remote ($N = 47$); blue: non-remote ($N = 47$). Line: WLS fit.

non-parents based on age, education, race, and state. Occupations are assigned at the time of observation. We measure four outcomes: employment last year, log usual hours, log earnings, and log hourly wages. We partition the pre-COVID years into four five-year periods (2001–05, 2006–10, 2011–15, and 2016–19) and define 2022–24 as the post period, with 2016–19 as the reference. Five-year bins provide enough observations for precise occupation-level estimation, and four pre-treatment periods enable documenting child penalty trajectories.

CPS. We complement the ACS with the CPS monthly files, 2000–2025, excluding individuals whose rotation group overlaps with 2020. The CPS rotating panel allows observing individuals before and after their first child's birth. Following Gulek (2024), individuals who have no child in their first interview cycle but have a newborn in their second cycle are assigned event time $t = -1$, providing a clean pre-birth observation without matching. Since the post-COVID window is short and occupation-level estimation requires statistical power, we rely on the ACS for our main results and use the CPS as a robustness check.

ATUS. We use the American Time Use Survey (ATUS), 2003–2024 (excluding 2020–21), to understand how remote work may differentially affect parents' productivity across genders. The ATUS is a one-day time diary administered to a subsample of CPS respondents (∼10,000 per year). It records *secondary childcare*: minutes during which a respondent supervises an own child while engaged in another activity. Occupations are classified as remote or non-remote using the same binary median split as in the main analysis.

3 Empirical Strategy

Our empirical design has three layers: we first show how we estimate child penalties across occupations and time periods, then document pre-trends in child penalties across remote and non-remote occupations, and finally apply a SC-DiD algorithm to adjust for these pre-trends and identify the causal effect of remote work intensity on child penalties.

3.1 Estimating child penalties

The event-study approach to estimating child penalties uses panel data on men and women who become parents. For each occupation o and gender g, we estimate:

$$\begin{aligned} Y^g_{iot} &= \sum_{j\neq -1} \alpha^g_{o,j}\, \Delta D_{i,t-j} + \theta^g_o W_{it} + \epsilon^g_{iot}, \\ \ln(Y^{g,o}_{it}) &= \sum_{j\neq -1} \alpha^g_{o,j}\, \Delta D_{i,t-j} + \theta^{g,o} W_{it} + \epsilon^{g,o}_{it}, \end{aligned} \tag{1}$$

where $\Delta D_{i,t} = 1$ if individual i had their first child at time t, and j indexes event time relative to first birth, with $j = -1$ as the omitted category. The first equation estimates employment on the full sample; the second estimates conditional outcomes (hours, income, wages) among the employed, hence occupation appears as a superscript denoting sample selection rather than a subscript. The controls W_{it} include age and calendar year fixed effects (Kleven et al., 2019, 2024a), and education fixed effects following Cortés and Pan (2023). Education controls are particularly important in our setting: if the pandemic changed the timing of parenthood differentially for high- and low-skill workers, omitting education would amplify the bias for precisely the occupations where remote work expanded most.

This design historically required panel data, not because it exploits within-individual variation, but to observe eventual parents before they become parents. Two methods overcome this limitation when available panels are too small. Kleven (2025) developed the *pseudo-panel* approach, which uses exact matching to construct counterfactual non-parents from large cross-sections. Building on this work, Gulek (2024) showed that rotating panels can estimate (1) directly, because the regression requires only one pre-birth observation ($t = -1$) per individual, exactly what the CPS rotating panel provides. Since our setting requires statistical power for occupation-level estimation with a short post-treatment window, we rely on the ACS pseudo-panel for our main results and use the CPS rotating panel as a robustness check.

We implement the pseudo-panel as follows. In ACS year t, we observe parents whose eldest child is age $a \in \{0, 1, \ldots, 10\}$. Since our main outcomes refer to the prior year (when children were approximately one year younger), we assign event time $j = a - 1$ for parents with children aged $a \geq 1$. We do not assign $j = -1$ to parents with children aged 0, because imprecise birth timing within the survey year would contaminate the pre-birth observation. Instead, for $j = -1$, following Kleven (2025), we match parents to non-parents from the prior ACS year who are one year younger and share the same education, race, and state. We match on $k = 1$ rather than $k \in \{1, \ldots, 5\}$ as in Kleven (2025): the established absence of pre-trends makes additional matches unnecessary, and $k > 1$ would pull in the excluded 2020–21 ACS years.

For employment outcomes, we convert level estimates to percentage child penalties following Kleven et al. (2019): $\alpha^{\%} = \alpha/(\bar{Y}_{\text{parents}} - \alpha)$, where $\alpha < 0$ denotes a penalty and $\bar{Y}_{\text{parents}}$ is the mean outcome for parents. Log outcomes are already in percentage terms.

For occupation-level estimation, small cell sizes make the fully dynamic specification imprecise. We therefore also estimate a static version with a single parenthood indicator D_{it}, yielding one precisely estimated parameter per occupation-period cell. The pre-trend analysis in Section 3.2 uses this static version; the SC-DiD stage returns to child-age heterogeneity.

3.2 Child penalty trajectories across remote and non-remote occupations

To document how child penalties have evolved over time across remote and non-remote occupations, we estimate a version of (1) that separates the baseline child penalty from its period-specific deviations. For each occupation group o and gender g:

$$Y^g_{iot} = \gamma^g_o D_{it} + \sum_{\tau \neq \tau_0} \alpha^g_{o,\tau} \left(D_{it} \times \mathbb{1}\{t \in \tau\}\right) + \theta^g_o W_{it} + \epsilon^g_{iot}, \tag{2}$$

where γ^g_o captures the baseline child penalty for gender g and occupation group o that is common across time, and $\alpha^g_{o,\tau}$ captures the differential child penalty in period τ relative to the reference period τ_0 (2016–19). An analogous specification with $\ln(Y^{g,o}_{it})$ is estimated for conditional outcomes, as in (1). We estimate (2) separately for remote ($\Delta R_o \geq$ median) and non-remote ($\Delta R_o <$ median) occupations.

Figure 2: Child Penalty Trajectories: Remote vs. Non-Remote (ACS, Employment)

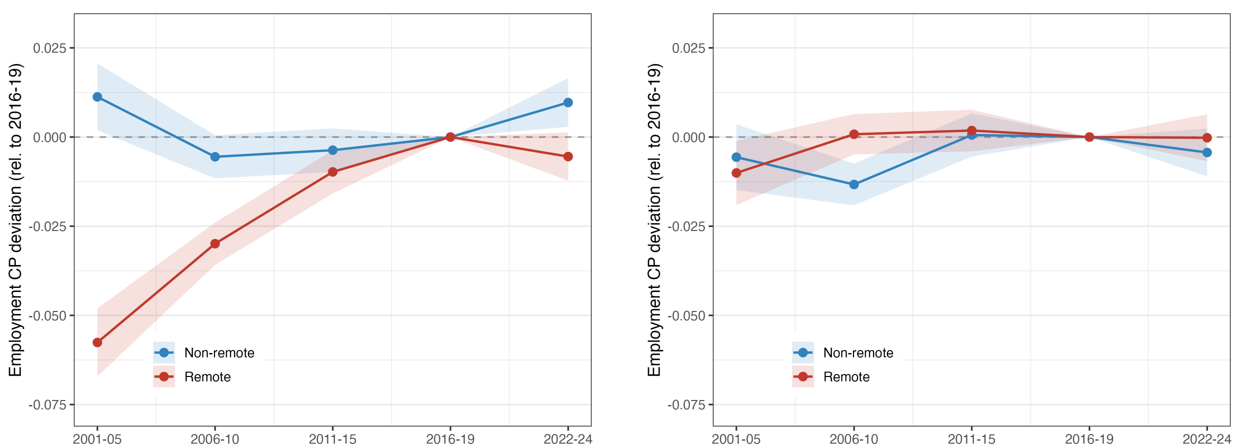

(a) Women (b) Men

Employment child penalty deviations $\alpha^g_{o,\tau}$ from (2), expressed relative to the reference period (2016–2019, normalized to zero). Red: remote occupations ($\Delta R_o \geq$ median). Blue: non-remote. ACS pseudo-panel, robust SEs. Women (left): the penalty coefficient in remote occupations trends upward before COVID, indicating the child penalty was shrinking. Men (right): no differential pattern. Appendix Figure A.1 shows all four outcomes by gender.

3.3 Adjusting for differential trends using synthetic controls

We model the occupation-level child penalty as:

$$\alpha^g_{o,\tau} = \beta^g \Delta R_o \cdot \mathbb{1}\{\tau = \tau_1\} + f_o + f_\tau + \lambda_o' \theta_\tau + \epsilon_{o,\tau}, \tag{3}$$

where ΔR_o is the change in remote work share induced by the pandemic, f_o and f_τ are occupation and time-period fixed effects, and $\lambda_o' \theta_\tau$ captures unobserved confounders generating differential trends. Since we cannot observe these confounders, we use synthetic controls to proxy for them. Identification requires that the interactive fixed effects $\lambda_o'\theta_\tau$ are low-rank relative to the number of pre-periods and matching moments, so that the synthetic control weights can recover the unobserved loadings from the observed pre-treatment trajectories. Put differently, we require that the residual variation in remote work adoption between occupations and their synthetic controls is orthogonal to occupation-specific shocks affecting child penalties in the post period.

The empirical challenge is that ΔR_o is not binary: some occupations are far more exposed than others, and we want to exploit this variation for identification. The SDiD of Arkhangelsky et al. (2021) does not apply, as it requires a binary treatment. We use the Synthetic IV algorithm of Gulek and Vives-i Bastida (2024), which converges to the following estimator when the instrument equals the treatment; that is, we assume no endogeneity and are concerned only with omitted variable bias. The algorithm reduces to: (i) construct synthetic controls matching pre-treatment trajectories, (ii) debias both outcomes and treatment using the same weights, (iii) regress debiased outcome changes on debiased treatment.

Intuitively, we want to create a synthetic occupation that would have followed the same child penalty trajectory absent the remote work shock. We find synthetic control weights, separately by gender, that best match pre-COVID child penalty trends jointly across employment and hours. Joint matching improves the signal-to-noise ratio when outcomes are driven by related confounders (Sun et al., 2025), ensures an apples-to-apples comparison (the same weights are used for all outcomes, so differences in treatment effects cannot be driven by differences in weights), and provides an additional diagnostic: SC-DiD should not create pre-trends in untargeted outcomes. The hours coefficient is similar under hours-only separate weights ($+0.0019$, $p = 0.042$), confirming that the result does not depend on joint matching. Since we have many donor units and few pre-periods, we further enforce the weights to match trajectories across four child-age groups: age 0, ages 1–2, 3–5, and 6–10.[4]

Denote these four age groups by $j \in \{j_1, j_2, j_3, j_4\}$ and the five-year time periods by $\tau \in \{\tau_{-3}, \tau_{-2}, \tau_{-1}, \tau_0, \tau_1\}$. The algorithm proceeds as follows:

Algorithm.

[4] Of the 94 occupations in the data, 17 have insufficient cell sizes to compute percentage-normalized employment child penalties across all child-age bins and time periods. The SC-DiD analysis uses the remaining 77 occupations.

1. **Estimate CPs.** Run (1) per occupation o, child age j, and period τ, yielding $\{\alpha_{o,\tau}^{g,j}\}$.

2. **Residualize.** In pre-periods ($\tau \leq \tau_0$), remove occupation and period fixed effects for each gender and child-age group: $\alpha_{o,\tau}^{g,j,\text{res}} = \alpha_{o,\tau}^{g,j} - \bar{\alpha}_o^{g,j} - \bar{\alpha}_\tau^{g,j} + \bar{\alpha}^{g,j}$.

3. **SC weights.** For each occupation and gender, find non-negative weights summing to one that minimize:
$$\min_{w_o^g} \sum_{\tau \leq \tau_0} \sum_j \sum_{y \in \{\text{emp,hrs}\}} \left(\alpha_{o,\tau}^{y,g,j,\text{res}} - \sum_{o'} w_{o,o'}^g \, \alpha_{o',\tau}^{y,g,j,\text{res}} \right)^2, \tag{4}$$
where $w_{o,o'}^g$ is the weight that occupation o' receives as a synthetic control for occupation o, with $w_{o,o}^g = 0$ (each occupation is excluded from its own donor pool), and the summation over y enforces joint matching across employment and hours outcomes.

4. **Debias.** Using the same weights, construct debiased versions of both the static child penalty (averaged across child-age groups) and the dynamic child penalties (for each child-age group separately): $\tilde{\alpha}_{o,\tau}^g = \alpha_{o,\tau}^g - \sum_{o'} w_{o,o'}^g \alpha_{o',\tau}^g$ and $\tilde{\alpha}_{o,\tau}^{g,j} = \alpha_{o,\tau}^{g,j} - \sum_{o'} w_{o,o'}^g \alpha_{o',\tau}^{g,j}$. Using a single set of weights for both static and dynamic effects ensures a consistent comparison. We also debias the treatment: $\tilde{R}_o^g = \Delta R_o - \sum_{o'} w_{o,o'}^g \Delta R_{o'}$. This step is non-standard: in the classic synthetic control framework, donor units are never treated, so $\tilde{R}_o^g = \Delta R_o$. Our generalization allows all occupations to serve as donors even when all are treated to varying degrees.[5]

5. **Estimate.** Regress the post-minus-pre change in debiased CPs on the debiased treatment, using only the post period:
$$\begin{aligned} \text{static:} \quad & \Delta\tilde{\alpha}_o^g = \beta^g \tilde{R}_o^g + \epsilon_o^g, \\ \text{dynamic:} \quad & \Delta\tilde{\alpha}_o^{g,j} = \beta^{g,j} \tilde{R}_o^g + \epsilon_o^{g,j}. \end{aligned} \tag{5}$$

In practice, the static estimand $\Delta\tilde{\alpha}_o^g$ is computed as the equally weighted average of child-age-bin-specific debiased changes: $\Delta\tilde{\alpha}_o^g = \frac{1}{J}\sum_j [(\tilde{\alpha}_{o,\tau_1}^{g,j} - \tilde{\alpha}_{o,\tau_0}^{g,j})]$.

First, for inference, we use randomization inference by permuting the treatment ΔR_o across occupations 1,000 times, keeping SC weights fixed.[6] We compute two-sided p-values as the fraction

[5] The debiased treatment retains substantial variation: for women, $\text{SD}(\tilde{R}_o^g) = 6.7$ compared to $\text{SD}(\Delta R_o) = 6.9$, with a correlation of 0.93 between the raw and debiased treatment. Appendix Figure A.9 shows the scatter.

[6] Fixing SC weights during randomization inference is standard: the weights are functions of the outcome, not the

of permuted test statistics that exceed the actual estimate in absolute value. For confidence intervals, we invert the p-value under a normal approximation.

Second, we also show robustness to using the dummy remote indicator in addition to the continuous treatment version. For the dummy specification, donors for each remote occupation are drawn from non-remote occupations, and vice versa.

Third, we weight regressions by occupation size; results are robust to inverse-variance weights.

4 Results

Figure 3 reports SC-DiD estimates for four outcomes: employment, and the natural logarithms of hours worked, income, and hourly wages. Our preferred specification uses four child-age bins $j \in \{0, 1\text{–}2, 3\text{–}5, 6\text{–}10\}$, joint matching across employment and hours, and size weights; we assess robustness across 12 specifications varying child-age bins, matching moments, and weighting schemes (Appendix Figure A.3). A positive coefficient indicates that the child penalty is becoming less negative, that is, the penalty is shrinking.

Our first result is that remote work does not reduce the employment child penalty for mothers. The point estimate in our preferred specification is negative (-0.0025, $p = 0.303$), indicating that if anything the employment penalty increases in occupations that adopted more remote work. This finding contrasts with the positive correlation between remote work and mothers' employment documented by Harrington and Kahn (2026) using pre-COVID data. As we show in Section 3.2, the employment penalty coefficient in remote occupations was already increasing throughout 2001–2019 (Figure 2); because the coefficient is negative at baseline, this upward trend means the child penalty was shrinking before COVID. Non-remote occupations showed no such trend. A naive estimator would conflate this pre-existing improvement with the treatment effect. Our synthetic control adjustment accounts for the differential pre-trend, and the apparent positive association disappears.

The null employment effect also addresses a compositional concern: any differential sorting of mothers across occupations after COVID would change the employment child penalty in remote occupations. The absence of such an effect makes it unlikely that the conditional results below are

treatment assignment. Re-computing weights for each permutation would conflate the SC step with the inference step.

Figure 3: SC-DiD: Effect of Remote Work on Child Penalties

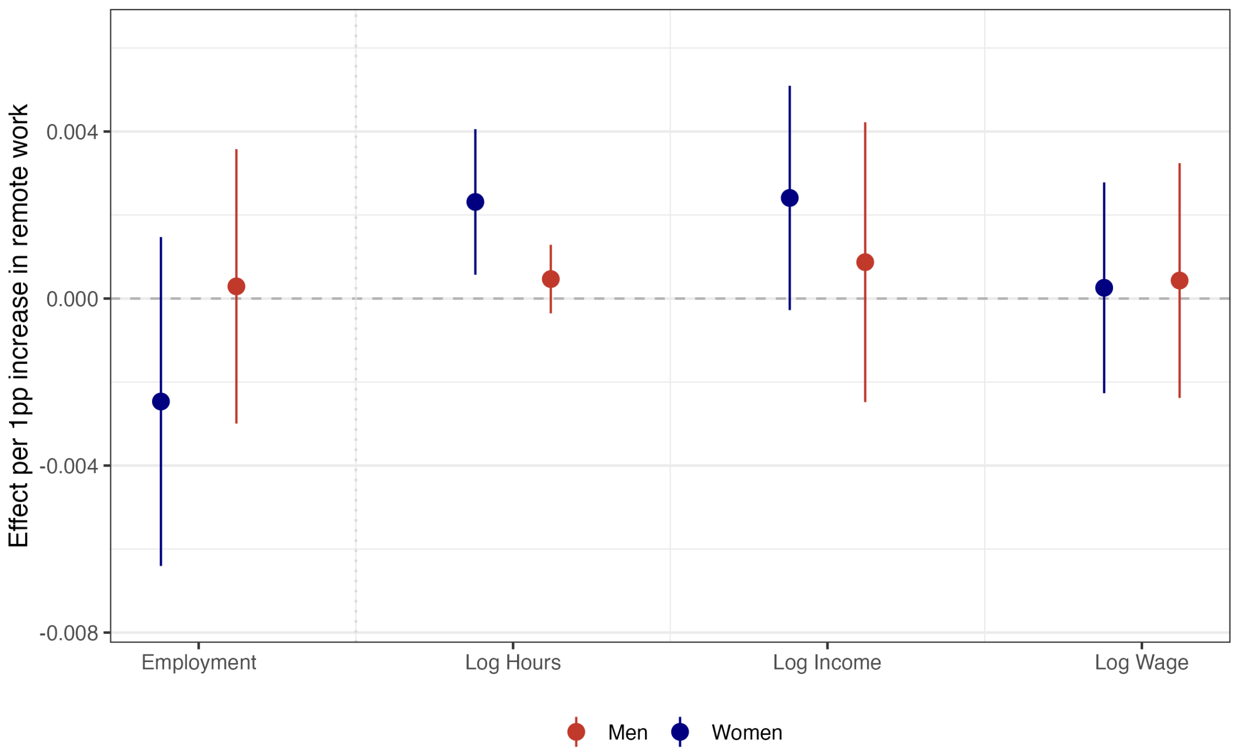

Notes: Each panel reports the SC-DiD treatment effect of occupation-level remote work adoption (ΔR_o, continuous, per 1 pp) on the child penalty for a given outcome. Blue markers denote women; red markers denote men. Bars show 95% confidence intervals from randomization inference (1,000 permutations). *Step 1.* For each occupation o, gender g, time period τ (2001–05, 2006–10, 2011–15, 2016–19, 2022–24), and child-age bin j (= 0, 1–2, 3–5, 6–10), we estimate the child penalty $\alpha_{o,j}^{g,\tau}$ from a regression of the outcome on parenthood indicators, controlling for age, calendar year, and education fixed effects. Employment penalties are estimated on the full sample interacting the dependent variable with an occupation indicator; conditional outcomes (hours, income, wages) are estimated within occupation among the employed. Employment effects are normalized by the counterfactual (predicted outcome absent the penalty). *Step 2.* For each occupation, we construct a synthetic control as a convex combination of all other occupations, with weights chosen to match pre-period residualized child penalties jointly across employment and hours and across all four child-age bins. We then compute the SC-debiased DiD estimand: $\widetilde{\Delta\alpha}_o^g = \frac{1}{J}\sum_j\left[(\tilde{\alpha}_{o,j}^{g,\text{post}} - \tilde{\alpha}_{o,j}^{g,\text{pre}})\right]$, where $\tilde{\alpha}$ denotes the SC-debiased penalty. The reported coefficient β is the slope from regressing $\widetilde{\Delta\alpha}_o^g$ on the SC-debiased treatment $\tilde{R}_o^g$, weighted by occupation size.

driven by selective sorting.

Second, we find robust evidence that remote work increases mothers' hours conditional on employment. In our preferred specification, each percentage point increase in an occupation's remote work share raises log hours by 0.23 percent ($p = 0.029$). The average remote occupation experienced an 8.9 pp larger increase in remote work share relative to the average non-remote occupation. The implied effect is $8.9 \times 0.23 \approx 2$ percent more hours worked. At baseline (2016–19), mothers in remote occupations worked 12.6 percent fewer hours than their pre-birth counterfactual, so this effect eliminates roughly one-sixth of the hours child penalty. For comparison, Ciasullo and Uccioli (2023) find that Australia's Fair Work Act, which entitled mothers to request reduced hours on permanent contracts, reduced the hours penalty by 22 percent. A market-driven change that eliminates 16 percent of the hours penalty is economically meaningful, particularly given that traditional family policies have had limited effects on child penalties (Kleven et al., 2024b). The hours effect is similar in magnitude across child-age bins, from newborns through age 10, indicating a persistent effect throughout the child-rearing years rather than one concentrated in infancy (Appendix Figure A.12).

Third, we find no statistically significant effect of remote work on mothers' income or hourly wages conditional on employment. The income point estimate is positive, similar in magnitude to the hours effect in the main specification, but not statistically significant (+0.0024, $p = 0.140$), remains insignificant and even changes sign across robustness checks. Wage point estimates are small in magnitude, mixed in sign across specifications, and all statistically insignificant (smallest $p = 0.115$, for a negative estimate). The significant finding is that remote work increases hours; we cannot reject null effects on income or wages. When hours increase, we would expect income to rise mechanically at the same wage and wages themselves to increase as mothers shift from part-time to full-time schedules. The fact that wages do not detectably increase despite the shift toward full-time work is consistent with the additional hours being less productive, as we document in Section 4.3, or with mothers accepting lower wages in exchange for the flexibility of remote work.

Fourth, for men, we find no average effects on any outcome; point estimates are close to zero and statistically insignificant across all four outcomes.

Together, these results indicate that remote work raises mothers' hours conditional on employment without translating into higher earnings or improved employment. This gendered contrast is consistent with remote work operating through caregiving flexibility rather than occupation-level demand shocks, which would affect both genders.

4.1 Robustness

We assess robustness along several dimensions. First, we vary the specification: across 12 combinations of child-age bins, matching moments, and weighting schemes, the hours effect is positive in all 12, statistically significant at the 5 percent level in five, and at the 10 percent level in nine (Appendix

Figure A.3). The employment effect is negative in all 12 specifications and marginally significant in one ($p = 0.055$). Income effects are insignificant throughout, and wage effects are indistinguishable from zero. These patterns confirm that the hours result is not an artifact of specification choice, and that the employment, income, and wage nulls are robust. A leave-one-out analysis confirms no single occupation drives the result: dropping each occupation and re-solving SC weights, the coefficient remains positive in all 77 cases (Appendix Figure A.13).

Overfitting of SC weights. We backtest the synthetic control weights following the best practice recommended by Abadie (2021). We restrict the training period to 2001–2015, assign a placebo treatment to 2016–2019, and test for spurious effects. All placebo tests for women yield precise nulls; the smallest p-value for women's log hours is 0.316. This suggests that the SC weights capture genuine treatment-period signal rather than overfitting to pre-period noise, though the test cannot rule out confounding from shocks concurrent with the pandemic.[7]

Sensitivity to treatment definition. Our main specification treats remote work intensity as continuous, exploiting variation in ΔR_o across occupations. An alternative is to dichotomize: classify occupations as "remote" or "non-remote" based on whether ΔR_o exceeds a threshold, and apply the Synthetic Difference-in-Differences (SDiD) estimator of Arkhangelsky et al. (2021). Dichotomizing requires choosing the percentile threshold defining treatment and how to aggregate child penalties across child-age bins. We vary the threshold from the median to the 80th percentile; Appendix Figure A.4 shows which occupations are treated at each threshold. Appendix Figure A.5 reports Wald-scaled SDiD estimates across these thresholds and three child-age bin specifications. The hours effect is positive across nearly all specifications, consistent with our main finding, though imprecise due to the loss of information from dichotomizing a continuous treatment. Income and wage effects are statistically insignificant throughout. The employment effect is sensitive to the threshold choice, turning positive at higher percentiles where fewer, more intensely treated occupations are classified as remote.

[7] Figure A.2 in the Online Appendix shows event-study estimates comparing OLS and SC-DiD with different training periods.

Alternative estimation sample. We replicate the analysis using the CPS rotating panel instead of the ACS pseudo-panel, following the rotating-panel approach of Gulek (2024). The rotating panel method does not rely on matching, but the CPS contains roughly nine times fewer parents with newborns per year. The CPS estimate is positive but imprecise; the two estimates cannot be statistically distinguished (Appendix Figure A.6).

Anticipatory occupational sorting. We test whether individuals anticipatorily sort into remote occupations before becoming parents. Using the pseudo-panel matching, we assign event time $t = -2$ to non-parents observed two years before their matched parent's first child and estimate whether future parents differentially enter remote occupations after COVID. Appendix Figure A.11 shows no evidence of sorting: the post-COVID coefficient is 0.002 for women ($p = 0.762$) and 0.008 for men ($p = 0.217$).

Heterogeneity by race and age. We explore heterogeneity along two dimensions (Appendix Figure A.10). The hours effect is positive for both white (+0.0023, $p = 0.022$) and non-white (+0.0017, $p = 0.251$) women; the two estimates are not statistically distinguishable. The effect is larger for young parents aged 25–34 (+0.0027, $p = 0.061$) than for older parents aged 35–45 (+0.0012, $p = 0.290$).

4.2 Distributional effects on hours

We investigate where in the hours distribution this effect originates. For each threshold $x \in \{2, 3, \ldots, 80\}$, we define $Y_x = \mathbf{1}\{\text{hours} \geq x\}$ conditional on employment and apply SC-DiD. The treatment effect is mechanically zero at the boundaries.

Figure 4 reports the treatment effect at each threshold for both genders. For women, the effect is near zero below around 13 hours per week, becomes positive and increases through 36 hours, where it peaks at +0.0022 ($p = 0.028$), and then declines back toward zero above 45 hours. The effect is statistically significant at the 5 percent level (pointwise) for thresholds between 17 and 25 hours and between 34 and 40 hours. Because each point measures the change in $\Pr(\text{hours} \geq x)$, a rising segment means probability mass is leaving those hours levels (more people now exceed each successive threshold), while a falling segment means mass is arriving. The pattern therefore implies that remote work shifts mothers from part-time hours (roughly 15–35 per week) toward full-time work (36–45 hours), with the largest effect at the conventional full-time threshold. This is consistent with the baseline hours penalty being largest at 35–40 hours (Appendix Figure A.7). Because the pseudo-panel does not track individuals, we cannot rule out that the distributional shift partly reflects changes in the composition of employed mothers rather than individual behavioral changes, though the null employment effect limits the scope for such compositional shifts.[8]

Figure 4: Distributional Effects of Remote Work on Hours

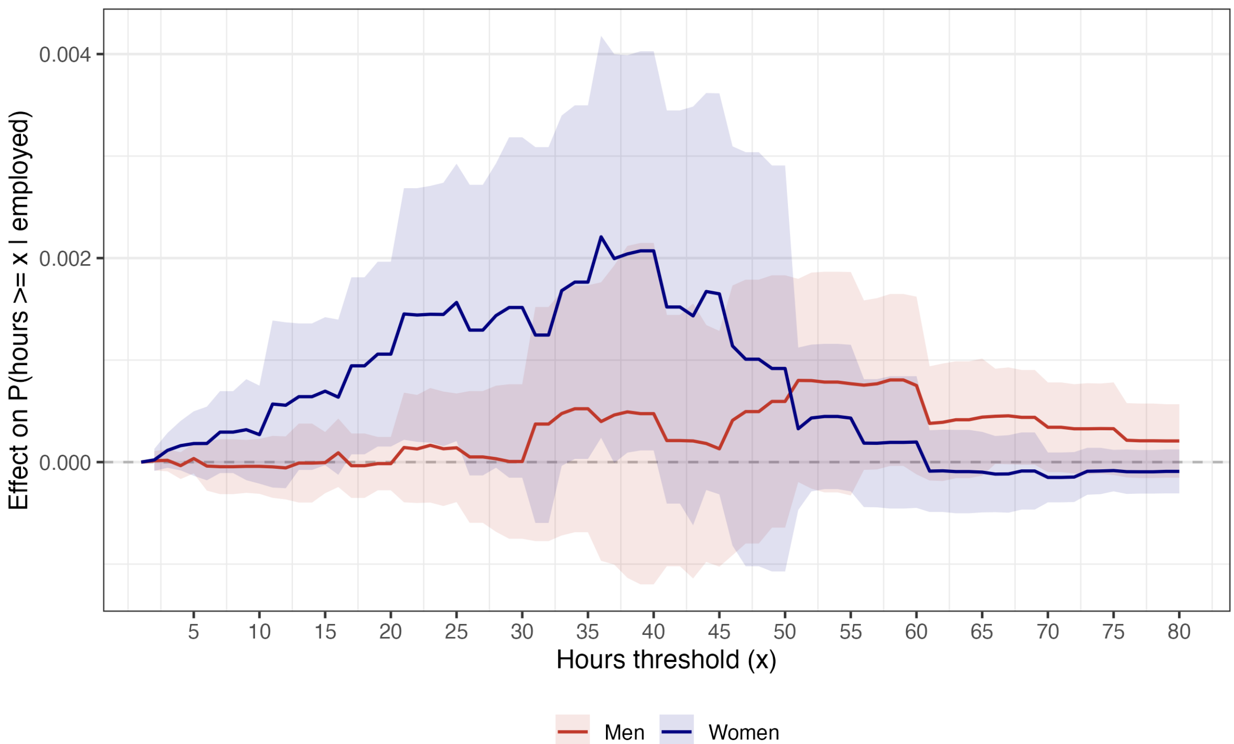

Notes: SC-DiD treatment effect of occupation-level remote work adoption (ΔR_o, continuous, per 1 pp) on Pr(hours $\geq x$ | employed) at each threshold $x = 2, 3, \ldots, 80$. Blue: women. Red: men. SC weights are the same as in Figure 3 (joint matching across employment and hours, four child-age bins, size-weighted). Shaded bands show 95% confidence intervals from randomization inference (1,000 permutations). The effect is mechanically zero at the boundaries. Pointwise p-values are unadjusted; the figure characterizes the shape of the distributional effect rather than providing 79 independent hypothesis tests.

For men, the treatment effect is near zero throughout, with a suggestive positive pattern above 50 hours that is imprecise.

[8]With large panels and individual fixed effects, the compositional concern would be eliminated entirely. In our pseudo-panel setting, it remains a potential concern.

4.3 Secondary childcare during work

Evidence shows hours increase without an apparent increase in income or wages. This is hard to reconcile since part-time positions come with wage penalties. Here we provide a novel mechanism that can help explain this pattern: remote work allows mothers to supervise children while working, lowering hourly productivity and absorbing the gains from additional hours.[9] We provide suggestive evidence on this channel using the ATUS, which records secondary childcare: minutes during which a respondent supervises a child while engaged in another activity. We focus on secondary childcare that occurs during paid work.

We collapse ATUS data to four cells (remote × period) using survey weights and estimate a difference-in-differences comparing remote and non-remote occupations before (2016–19) and after (2022–24) COVID. Pre-trends in secondary childcare are parallel for 2011–2019, justifying a simple DiD rather than SC-DiD. For inference, we permute the remote label across occupations (1,000 draws).[10]

Figure 5 presents the results. Panels (a) and (b) show mean secondary childcare during work by period, separately for fathers and mothers. In both panels, remote and non-remote occupations follow nearly identical trends through 2011–19, then diverge sharply in 2022–24. The divergence is larger for mothers: secondary childcare during work in remote occupations rises from 12 to 31 minutes per day for mothers, compared to 13 to 23 minutes for fathers, while non-remote occupations changed from 10.5 to 14.7 minutes for mothers and remained stable at 4.4 minutes for fathers. The difference-in-differences estimate is 15.5 minutes per day for mothers ($p = 0.005$) and 10.7 minutes for fathers ($p = 0.004$), relative to pre-treatment means of 11.6 minutes for mothers and 12.5 minutes for fathers, implying a 133 percent increase for mothers and an 86 percent increase for fathers.[11]

Panel (c) shows where in the distribution this effect originates. We estimate the DiD at each percentile. Below the 87th percentile, the treatment effect is zero for both genders: most parents

[9] Alternative mechanisms may coexist: mothers may also accept lower hourly wages in exchange for the flexibility of remote work (Mas and Pallais, 2017), which would independently generate an hours-income disconnect.

[10] Periods before 2011 are excluded because ATUS redesigned the secondary childcare module in 2010, creating a measurement break. The three periods used (2011–15, 2016–19, and 2022–24) provide one pre-trend check and one treatment comparison. Placebo tests comparing 2011–15 to 2016–19 yield near-zero estimates for both genders.

[11] Fathers' pre-treatment levels differ substantially between remote (12.5 min) and non-remote (4.4 min) occupations. A common proportional increase in childcare during work would generate a larger absolute DiD for the remote group even absent a treatment effect. This concern is mitigated for mothers, where pre-treatment means are closer (11.6 vs. 10.5 min).

Figure 5: Secondary Childcare During Paid Work (ATUS)

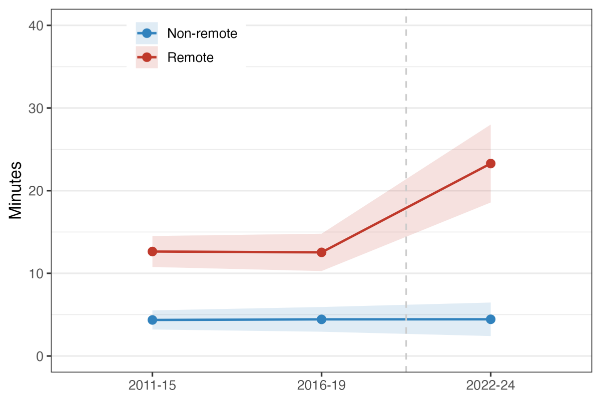

(a) Fathers

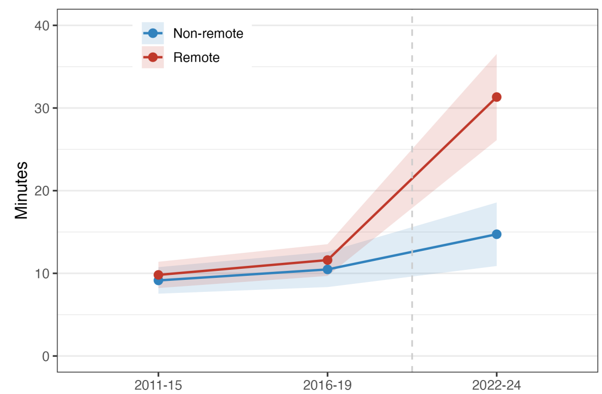

(b) Mothers

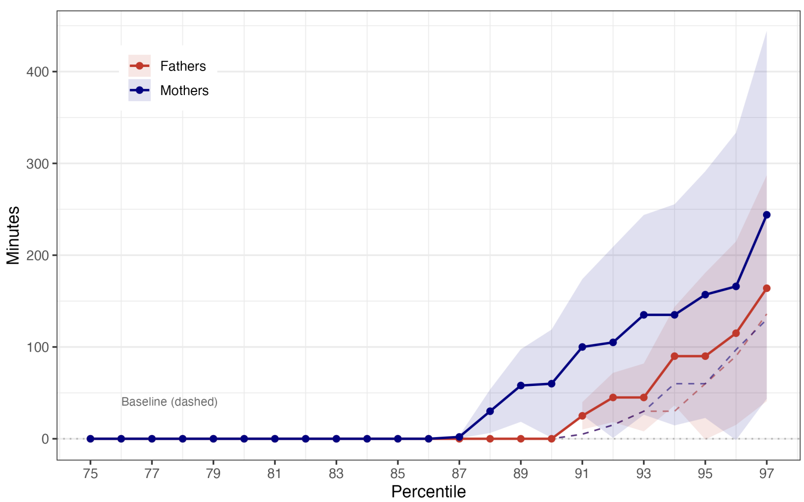

(c) Quantile treatment effects

Notes: Panels (a)–(b): Mean secondary childcare minutes during paid work, by period and remote status. ATUS, employed parents aged 18–64, survey-weighted, 95% CIs. Remote and non-remote defined by median split of ΔR_o (same classification as main analysis). Periods before 2011 excluded due to ATUS measurement redesign. Panel (c): Difference-in-differences at each percentile (2016–19 vs. 2022–24). Individual-level data collapsed to four cells (remote $\times$ period) using survey weights. Solid lines show DiD estimates; dashed lines show pre-treatment baseline levels in remote occupations. 95% CIs from randomization inference (1,000 permutations of the occupation-level remote label).

in remote occupations do no secondary childcare during work, as was the case before COVID.[12] The entire mean effect is driven by the upper tail. At the 91st percentile, mothers' treatment effect reaches 100 minutes ($p = 0.008$) and fathers' reaches 25 minutes ($p = 0.001$). Mothers' effect rises earlier and more steeply than fathers', diverging from zero around the 87th percentile

[12]The ATUS is a one-day diary; over a typical week, a larger fraction of mothers likely supervises children during some work hours.

compared to the 91st for fathers, consistent with mothers bearing a disproportionate share of childcare responsibilities when working from home.

These results are consistent with the productivity channel: mothers above the 90th percentile of secondary childcare average more than 60 minutes per day of concurrent childcare and paid work, which likely contributes to the absence of significant income and wage gains. Fathers also increase childcare during work, but less than mothers, consistent with remote work shifting some caregiving to fathers without fully equalizing the division of childcare. Consistent with this interpretation, Alipour (2025) finds that working from home shifts women's time from paid work to domestic work in Germany, while men's time use is largely unchanged, suggesting that the blurred boundary between paid and domestic work is a general feature of remote work for women.

5 Conclusion

Remote work reduces the hours child penalty for mothers but does not significantly affect employment, income, or wage penalties. For the average remote occupation, the implied effect is approximately 2 percent more hours worked, eliminating one-sixth of the baseline hours penalty. We find no average effects on men. ATUS evidence suggests that the muted income and wage gains may reflect lower productivity: mothers in remote occupations sharply increased the time spent supervising children while working after COVID.

Three caveats apply, though our analysis addresses each to the extent possible. First, the ACS pseudo-panel does not track individuals, so we cannot fully separate within-individual behavioral changes from compositional shifts. However, the null employment effect limits the scope for such shifts. Second, Lightcast measures posted remote work availability rather than realized arrangements; if measurement error is classical, our estimates are attenuated and the true effects would be larger. Third, we cannot rule out that pandemic-concurrent shocks correlated with both remote work adoption and child penalty changes.

Occupations that eventually adopted remote work were already on differential child penalty trajectories before COVID, and any causal analysis must account for these pre-existing trends. Our combination of pseudo-panels and synthetic controls for continuous treatments addresses these trajectories, while excluding 2020–2021 avoids pandemic-induced structural breaks. Researchers can apply this approach wherever large cross-sectional data exist (Kleven et al., 2024a).

Remote work did not decrease mothers' employment penalties, but employment penalties in eventually-remote occupations did decrease throughout the 2000s and 2010s, well before COVID. Possible drivers include improvements in communication technology, online conferencing, and changing workplace norms, though the underlying mechanisms remain unidentified. Early cross-country evidence suggests heterogeneous effects: remote work reduces employment penalties in Latin America (Zarate, 2025), hours penalties in the United States, and neither in Sweden (Scott and Sundberg, 2025). Why effects differ across settings, whether due to childcare institutions, gender norms, or the nature of remote work itself, is an important question for future research.

References

Abadie, Alberto, "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects," *Journal of Economic Literature*, 2021, *59* (2), 391–425.

Aksoy, Cevat Giray, Nicholas Bloom, Steven J. Davis, Victoria Marino, and Cem Ozguzel, "Remote Work, Employee Mix, and Performance," Working Paper 33851, National Bureau of Economic Research 2025.

Alipour, Jean-Victor, "Does Remote Work Reinforce Gender Gaps in (Un)Paid Labor?," Rationality and Competition Discussion Paper Series 542, CRC TRR 190 Rationality and Competition 2025.

Arkhangelsky, Dmitry, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager, "Synthetic Difference-in-Differences," *American Economic Review*, 2021, *111* (12), 4088–4118.

Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying, "Does Working from Home Work? Evidence from a Chinese Experiment," *Quarterly Journal of Economics*, 2015, *130* (1), 165–218.

_ , Ruobing Han, and James Liang, "How Hybrid Working from Home Works Out," Working Paper 30292, National Bureau of Economic Research 2022.

Ciasullo, Ludovica and Martina Uccioli, "What Works for Working Mothers? A Regular Schedule Lowers the Child Penalty," 2023. Available at SSRN: `https://ssrn.com/abstract=4572399`.

Cortés, Patricia and Jessica Pan, "Children and the Remaining Gender Gaps in the Labor Market," *Journal of Economic Literature*, 2023, *61* (4), 1359–1409.

Emanuel, Natalia and Emma Harrington, "Working Remotely? Selection, Treatment, and the Market for Remote Work," *American Economic Journal: Applied Economics*, 2024, *16* (4), 528–559.

Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 2014, *104* (4), 1091–1119.

—, "Nobel Lecture: An Evolving Economic Force," *American Economic Review*, 2024, *114* (6), 1515–1539.

— **and Lawrence F Katz**, "A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-friendly Occupation," *Journal of Labor Economics*, 2016, *34* (3), 705–746.

Gulek, Ahmet, "Occupational Child Penalties," 2024. Available at: https://ahmetgulek.github.io/Gulek_CP_Occupations.pdf.

— **and Jaume Vives i Bastida**, "Synthetic IV Estimation in Panels," 2024. Available at: https://economics.mit.edu/sites/default/files/inline-files/Synthetic_Wald%20%283%29.pdf.

Hansen, Stephen, Peter John Lambert, Nicholas Bloom, Steven J Davis, Raffaella Sadun, and Bledi Taska, "Remote Work across Jobs, Companies, and Space," Working Paper 31007, National Bureau of Economic Research 2023.

Harrington, Emma and Matthew E Kahn, "Has the Rise of Work from Home Reduced the Motherhood Penalty in the Labor Market?," *National Tax Journal*, 2026, *79* (2).

Ho, Lisa, Suhani Jalota, and Anahita Karandikar, "Bringing Work Home: Flexible Work Arrangements as Gateway Jobs for Women in West Bengal," *Quarterly Journal of Economics*, 2025. Conditionally accepted.

Kleven, Henrik, "The Geography of Child Penalties and Gender Norms: Evidence from the United States," 2025. NBER working paper: https://www.nber.org/papers/w30176.

—**, Camille Landais, and Gabriel Leite-Mariante**, "The Child Penalty Atlas," *Review of Economic Studies*, 2024, p. rdae104.

—**,** —**, and Jakob Egholt Søgaard**, "Children and Gender Inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 2019, *11* (4), 181–209.

__ , __ , **Johanna Posch, Andreas Steinhauer, and Josef Zweimüller**, "Do Family Policies Reduce Gender Inequality? Evidence from 60 Years of Policy Experimentation," *American Economic Journal: Economic Policy*, 2024, *16* (2), 110–149.

Mas, Alexandre and Amanda Pallais, "Valuing Alternative Work Arrangements," *American Economic Review*, 2017, *107* (12), 3722–59.

Olivetti, Claudia and Barbara Petrongolo, "The Economic Consequences of Family Policies: Lessons from a Century of Legislation in High-Income Countries," *Journal of Economic Perspectives*, 2017, *31* (1), 205–30.

Scott, Dana and Elin Sundberg, "Flexibility for Both Parents: Remote Work and the Evolution of Child Penalties," 2025.

Sun, Liyang, Eli Ben-Michael, and Avi Feller, "Using Multiple Outcomes to Improve the Synthetic Control Method," *Review of Economics and Statistics*, 2025, pp. 1–29.

Zarate, Pablo, "Remote Work and Child Penalties," *Available at SSRN 5360652*, 2025.

A Online Appendix

Contents

Figure A.1: Child Penalty Trajectories by Outcome (ACS, Both Genders)

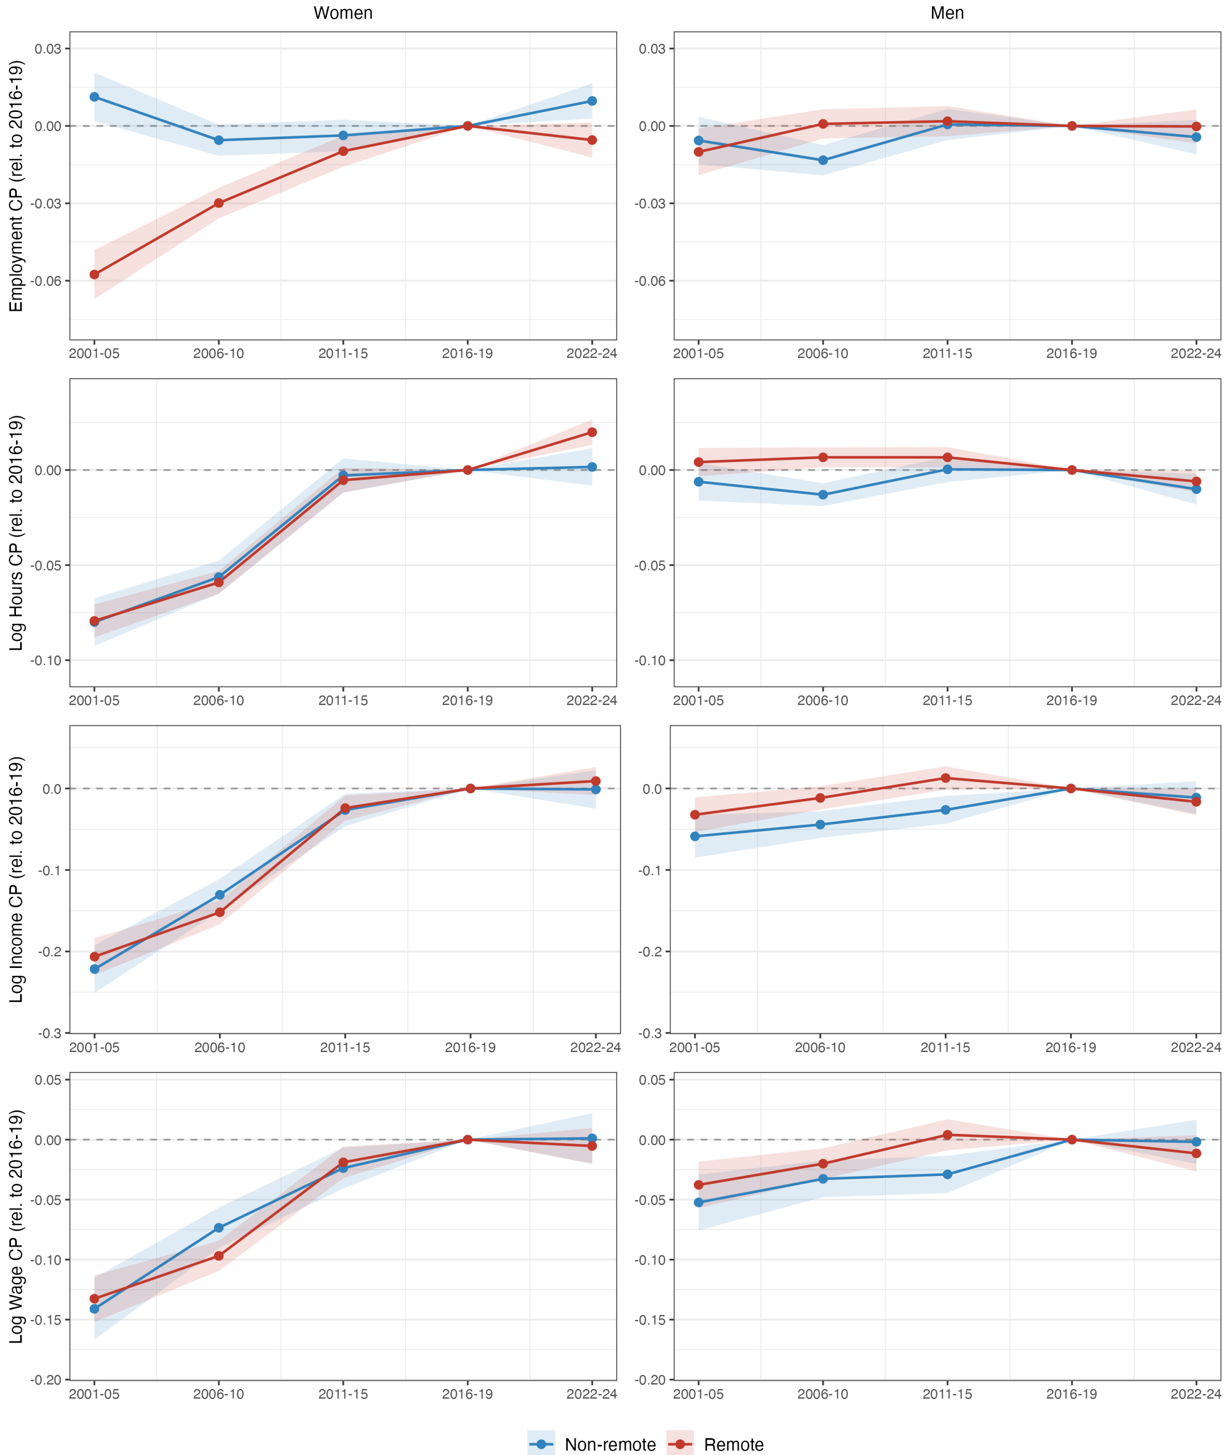

Notes: Child penalty deviations relative to the reference period (2016–19, normalized to zero). Left column: women. Right column: men. Red: remote occupations ($\Delta R_o \geq$ median). Blue: non-remote. ACS pseudo-panel, five periods: 2001–05, 2006–10, 2011–15, 2016–19 (reference), 2022–24. Y-axes are shared within each row for comparability across genders. Women's employment penalties show differential pre-trends across remote and non-remote occupations; conditional outcomes (log hours, log income, log wages) show parallel pre-trends, validating the parallel trends assumption for these outcomes.

Figure A.2: SC-DiD Period Event Study: Varying Training Windows (Women)

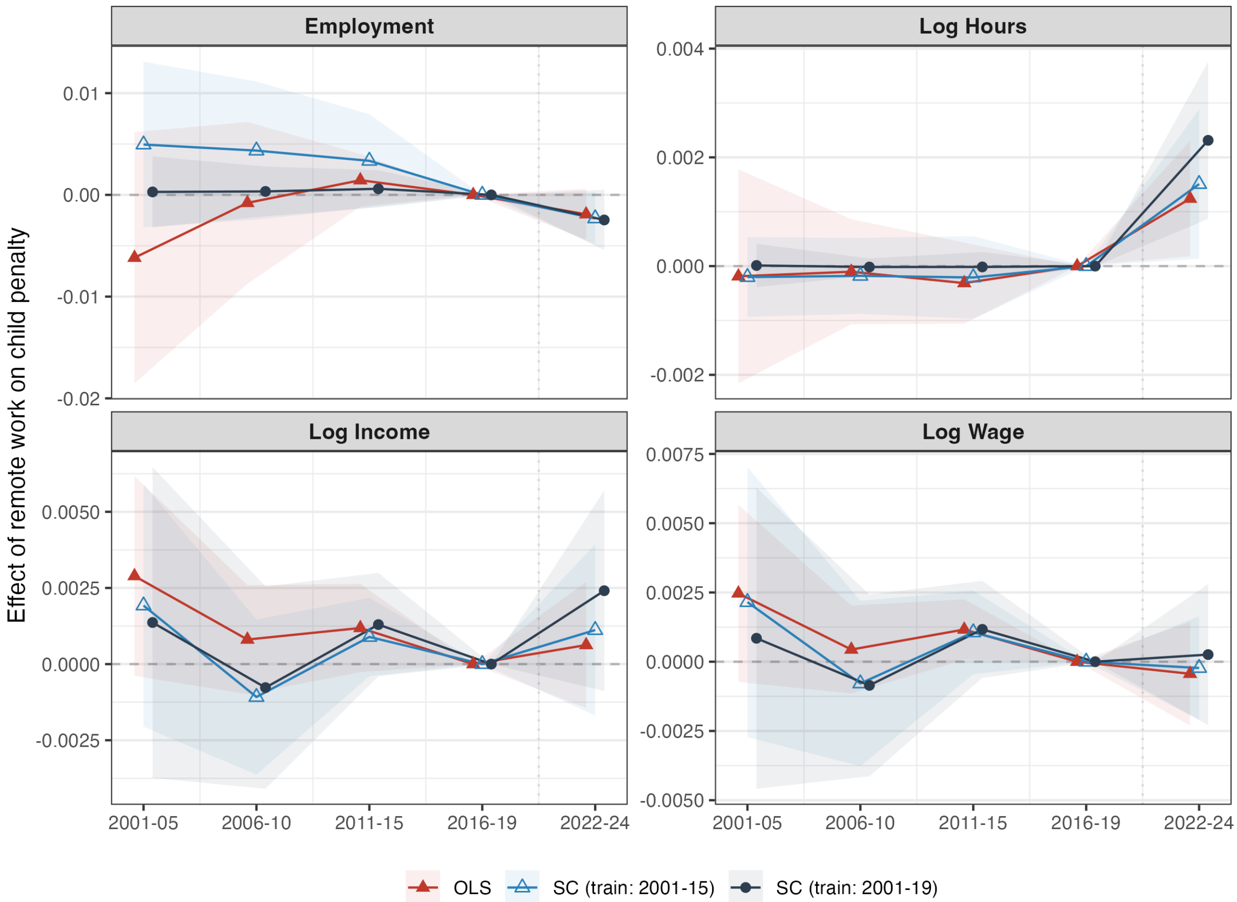

Notes: Period-by-period coefficients from regressing the (debiased) child penalty on (debiased) remote work intensity interacted with period indicators. Reference period: 2016–19. Three specifications: OLS (no debiasing, red triangles), SC-DiD trained on periods 2001–15 (blue triangles), and 2001–19 (dark circles, baseline). SC weights use joint matching across employment and hours (preferred specification: four child-age bins, size-weighted). Standard errors clustered at the occupation level. Women only. The progressive flattening of pre-period coefficients as the training window expands demonstrates that the synthetic control adjustment removes pre-existing trends, while the post-period coefficient remains stable across training windows.

Figure A.3: Robustness Across Specifications: Actual vs. Placebo (Women)

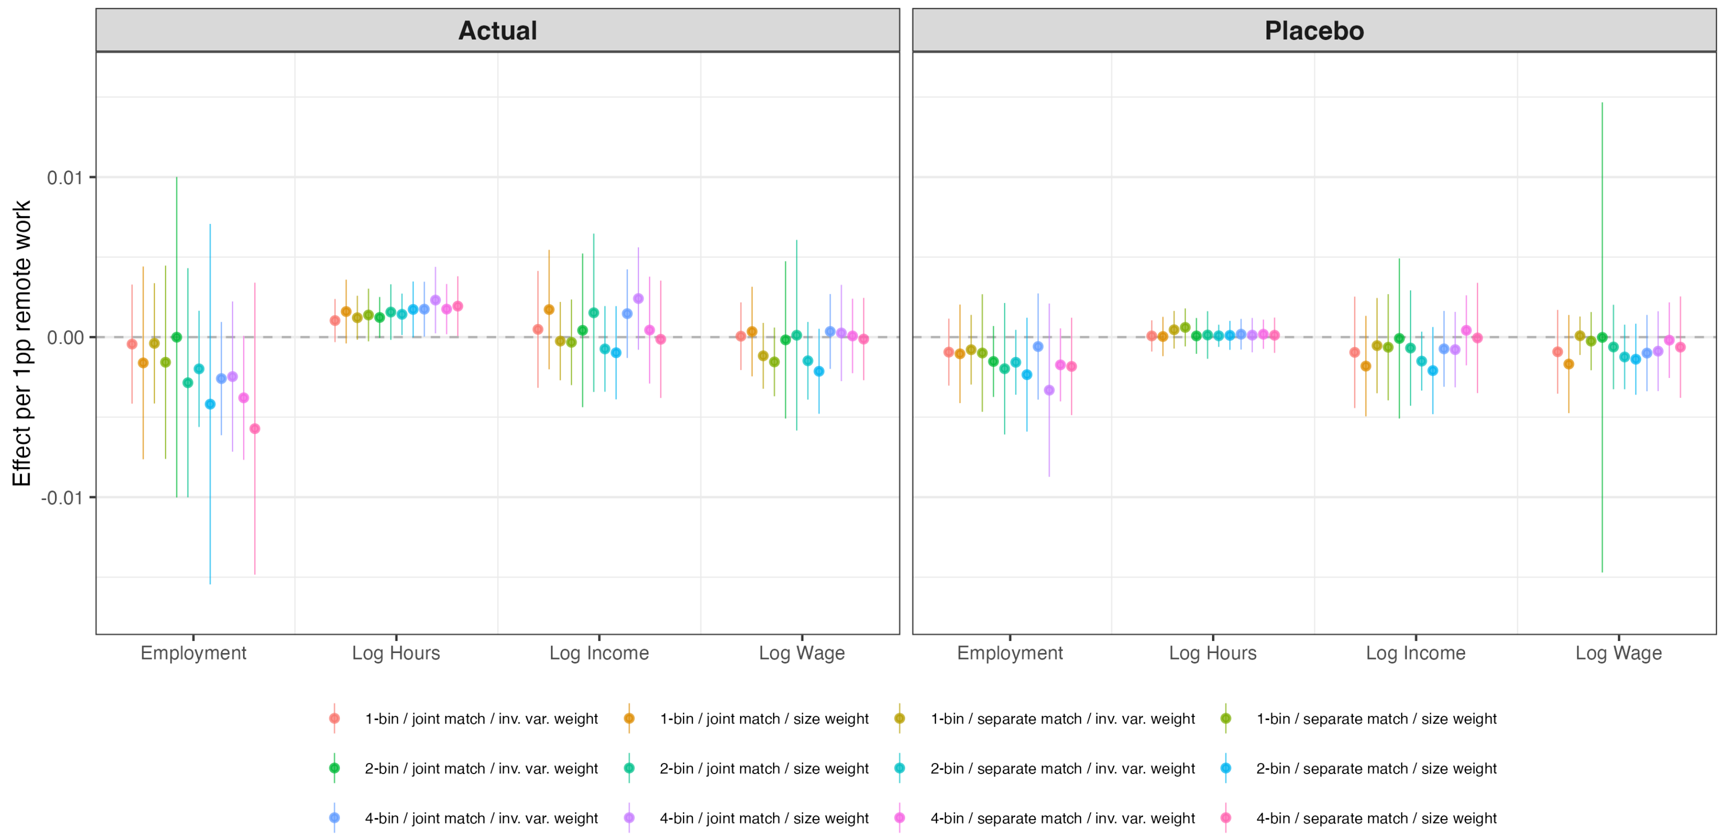

Notes: SC-DiD estimates across 12 specifications: 3 child-age bin configurations (1, 2, or 4 bins) × 2 matching approaches (separate weights per outcome vs. joint weights across employment and hours) × 2 weighting schemes (occupation size vs. inverse variance). Each point is one specification. Continuous treatment ΔR_o. 95% CIs from randomization inference (1,000 permutations). Left panel: actual estimates (2016–19 vs. 2022–24). Right panel: placebo backtest (training on 2001–15, testing at 2016–19). The hours effect is positive in all 12 actual specifications and zero in all 12 placebo specifications. The employment effect is negative in all 12 actual specifications; placebo employment estimates are near zero. Income and wages are null in both panels. For men, all actual and placebo estimates are close to zero across all four outcomes and all 12 specifications.

Figure A.4: Occupation Distribution with SDiD Treatment Thresholds

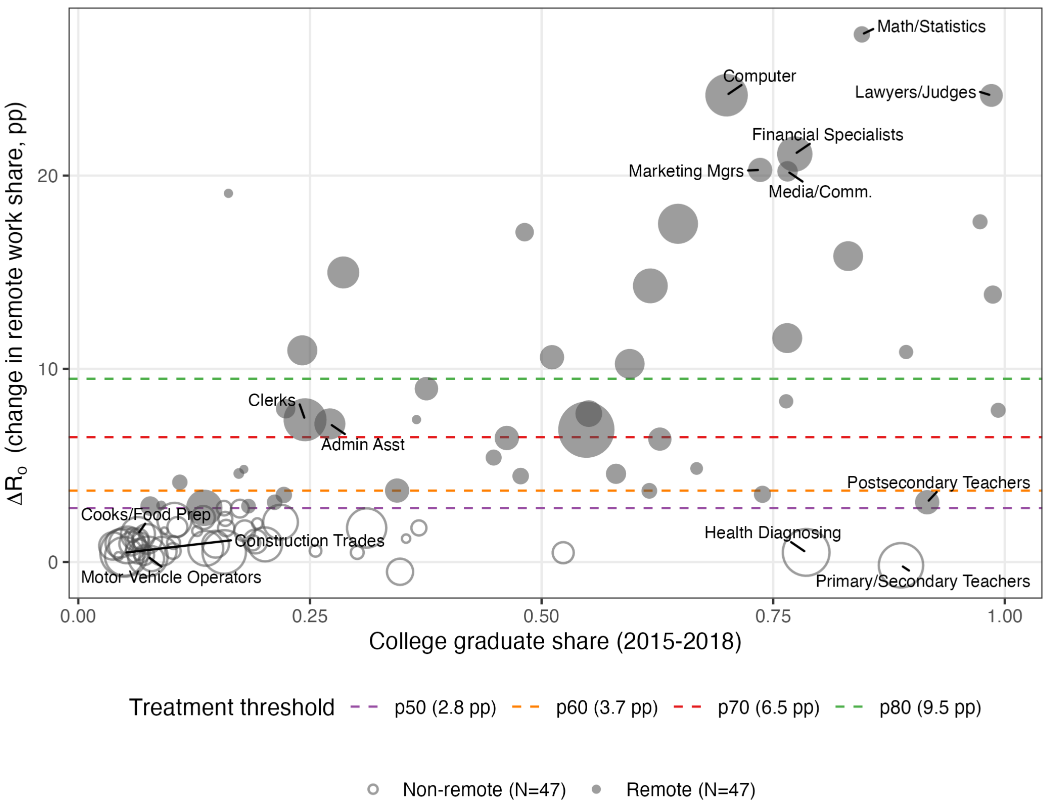

Notes: Each bubble is a 3-digit SOC occupation, sized by employment (ACS 2015–2018). x-axis: college graduate share. y-axis: ΔR_o, change in remote work share (pp). Horizontal dashed lines mark the treatment thresholds used in Figure A.5: at each threshold, occupations above the line are classified as treated. At the median (p50), half the occupations are treated; at p80, only the top 20% are treated. Filled circles: remote occupations (above median ΔR_o); open circles: non-remote.

Figure A.5: SDiD Robustness: Across Thresholds and Child-Age Bins (Women)

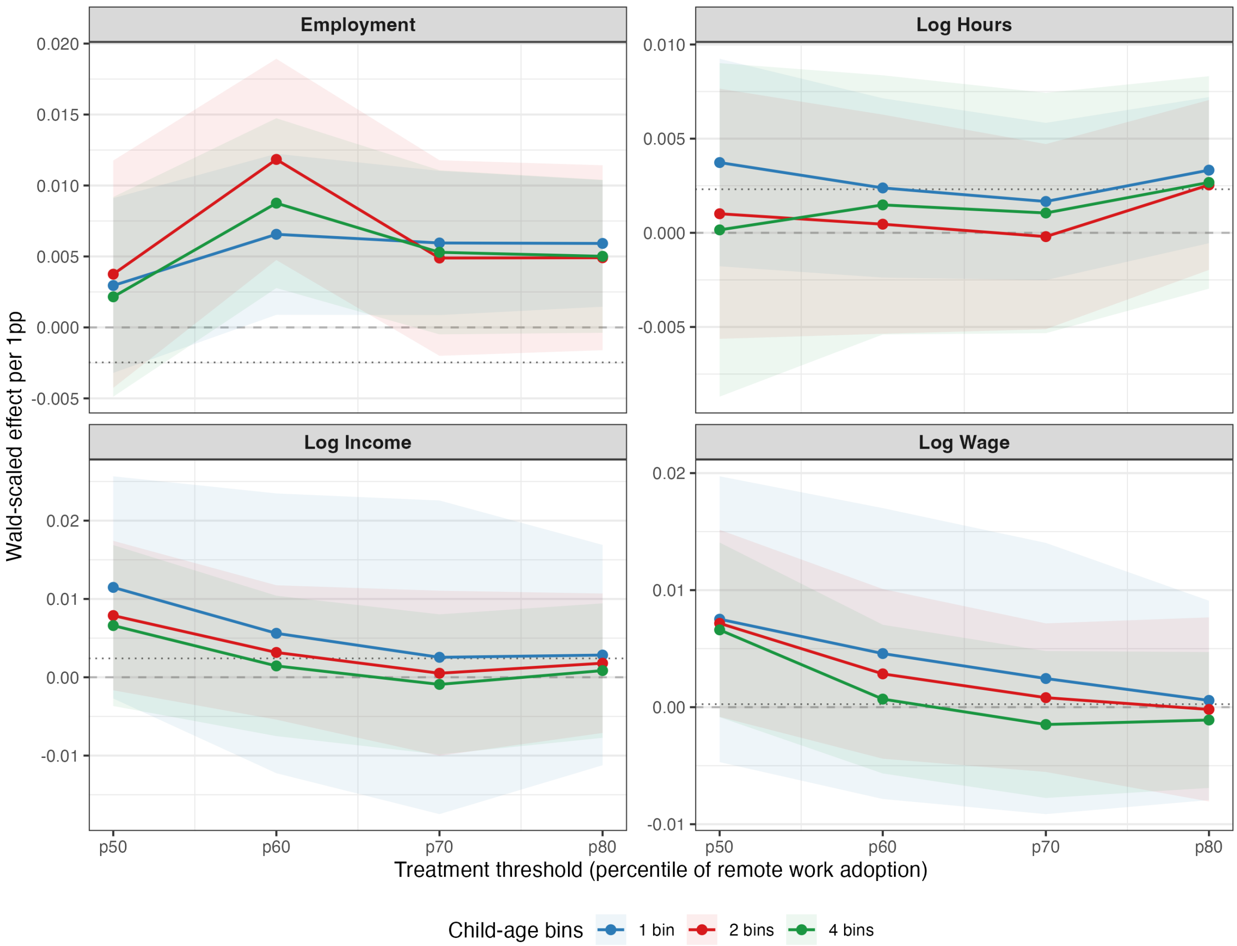

Notes: Synthetic Difference-in-Differences (SDiD) estimates of Arkhangelsky et al. (2021) for women, using a binary remote/non-remote classification. Each panel shows one outcome. The x-axis varies the percentile threshold that defines treatment from the median (p50) to the 80th percentile (p80): at the p-th percentile, occupations with ΔR_o above the p-th percentile are classified as treated (see Figure A.4 for which occupations fall above each threshold). As the threshold rises, the treated group narrows to the most intensely remote occupations. Three child-age bin specifications are shown (1, 2, or 4 bins). Coefficients are Wald-scaled by the mean gap in ΔR_o between treated and control groups at each threshold, making magnitudes comparable to the continuous SC-DiD in Figure 3. Dotted horizontal line: main SC-DiD continuous estimate ($\hat{\beta}$, es4jbin, joint, size-weighted). 90% CIs from randomization inference (1,000 permutations). The hours effect is positive across nearly all specifications, consistent with the main finding. Employment is sensitive to the threshold choice. Income and wage effects are statistically insignificant throughout. All SDiD estimates for men are null (not shown).

Figure A.6: Robustness: Pseudo-Panel (ACS) vs. Rotating-Panel (CPS)

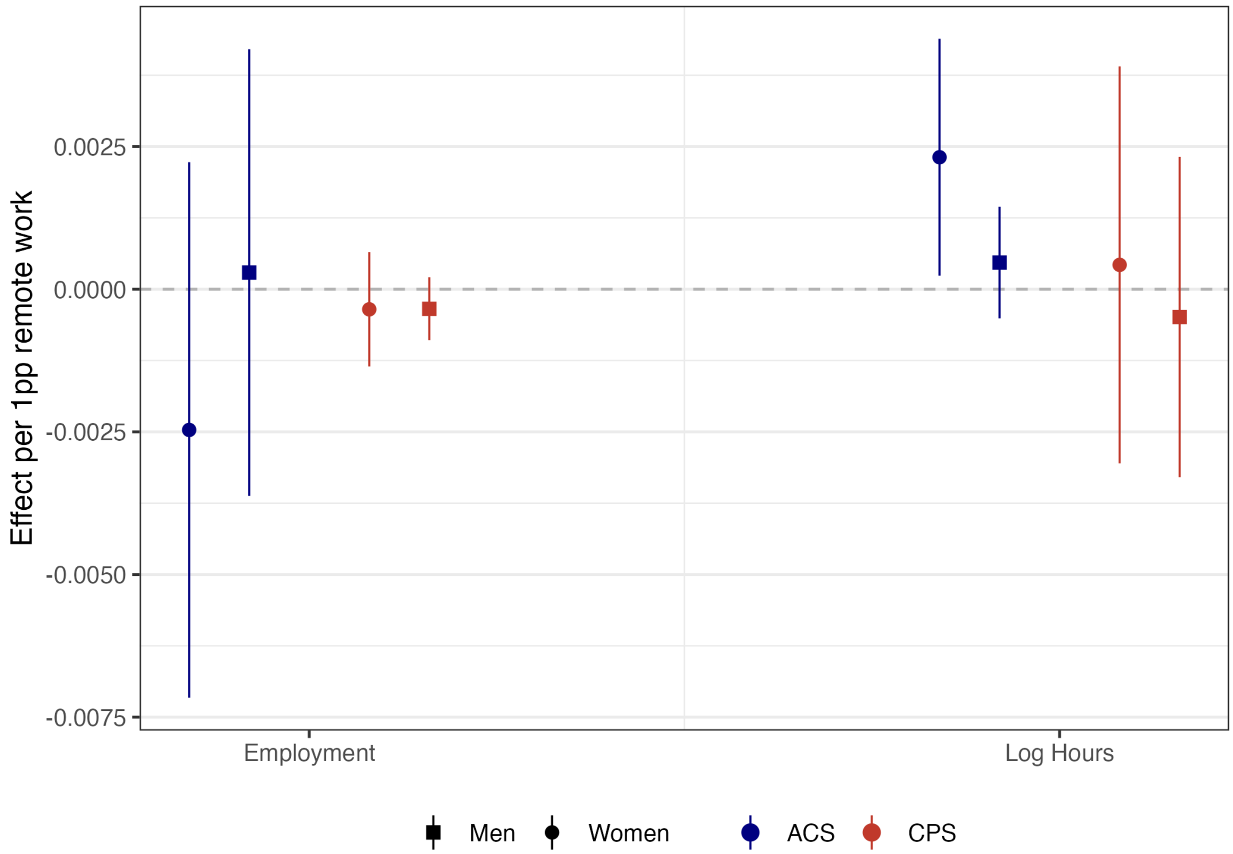

Notes: SC-DiD estimates comparing the pseudo-panel approach (ACS, blue) with the rotating-panel approach (CPS, red). The ACS pseudo-panel matches parents to observationally similar non-parents using age, education, race, and state; the CPS rotating panel directly observes individuals before and after their first child's birth without matching. Circles: women. Squares: men. Continuous treatment ΔR_o, four child-age bins, size-weighted. RI 95% CIs from 1,000 permutations. The CPS estimate is consistent with the ACS in sign and magnitude, though too imprecise to independently confirm the effect. The CPS contains $\sim 9\times$ fewer parents with newborns per year than the ACS, explaining the wider confidence intervals.

Figure A.7: Baseline Distributional Hours Penalties

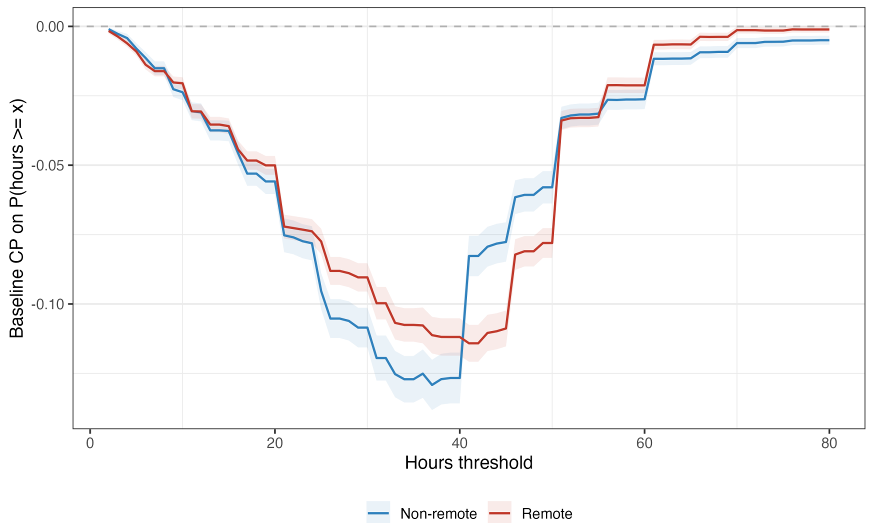

Notes: Baseline child penalty on Pr(hours $\geq x$ | employed) at each threshold x, estimated in the reference period (2016–19). Red: remote occupations. Blue: non-remote. The penalty measures how much motherhood reduces the probability of working at least x hours per week. Women's penalty is concentrated at $x = 35$–40 hours (the full-time threshold), with non-remote occupations exhibiting larger penalties than remote occupations at $x = 35$–45. This context helps interpret Figure 4: the treatment effect of remote work is largest where the baseline penalty is largest.

Figure A.8: Distributional Effects by Child Age (Women)

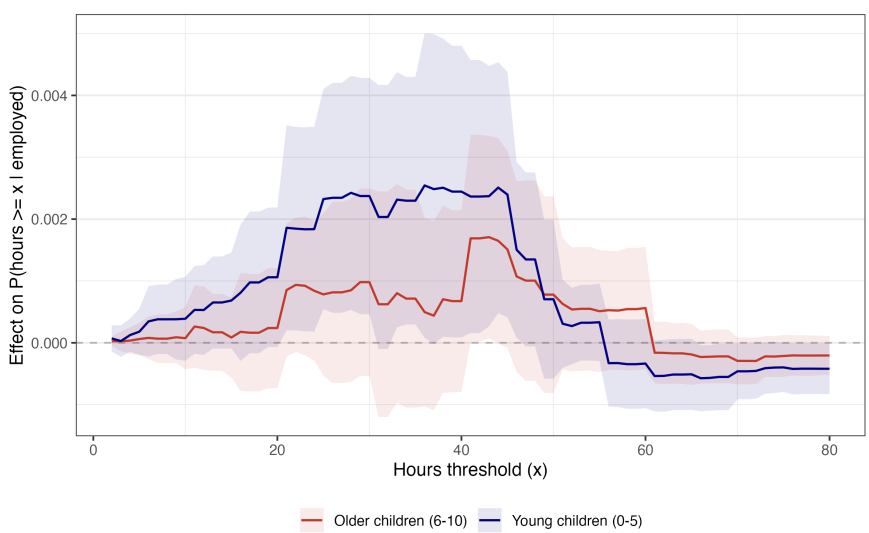

Notes: SC-DiD treatment effect of remote work on Pr(hours $\geq x$ | employed) at each threshold x, estimated separately for mothers with young children (age 0–5) and mothers with older children (age 6–10). Same SC weights as in the main specification (joint matching across employment and hours, four child-age bins, size-weighted). Continuous treatment ΔR_o. RI 95% CIs from 1,000 permutations. Point estimates are somewhat larger for mothers with young children, though the confidence intervals overlap at all thresholds. The mean effect on log hours is similar across child-age groups (Appendix Figure A.12).

Figure A.9: Raw vs. SC-Debiased Treatment (Women)

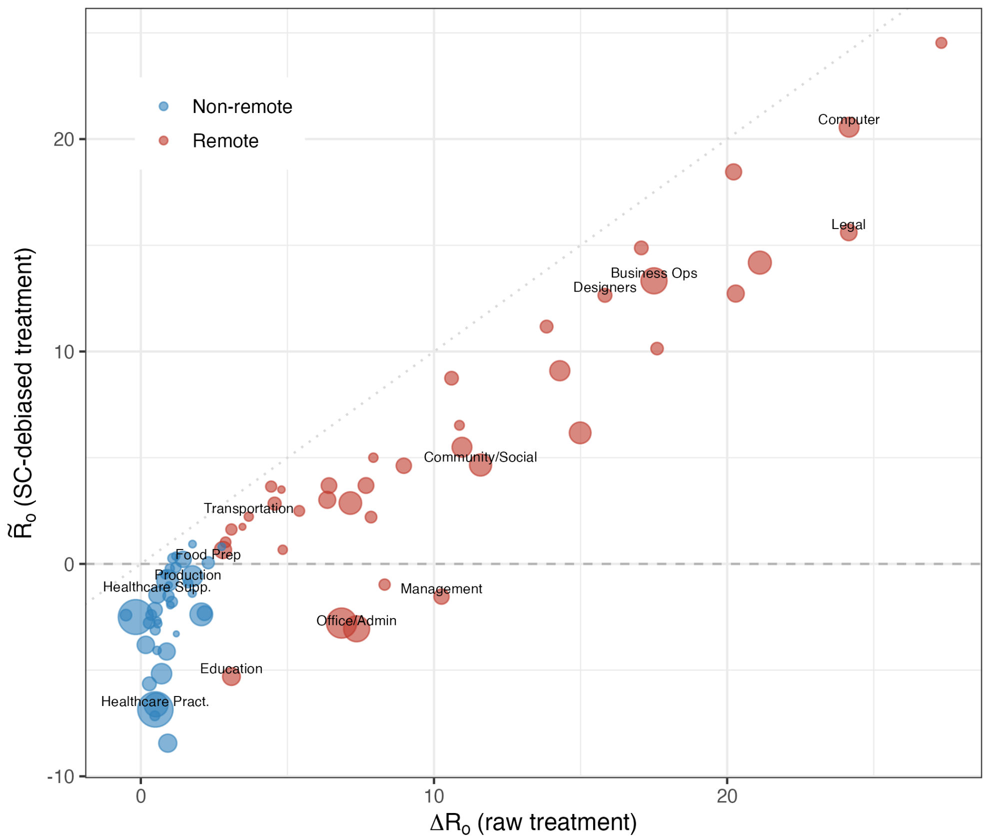

Notes: Each bubble is a 3-digit SOC occupation, sized by employment. x-axis: raw treatment ΔR_o (change in remote work share, pp). y-axis: SC-debiased treatment $\tilde{R}_o^g$ for women. Dotted line: 45-degree line. Dashed line: zero. The correlation between raw and debiased treatment is 0.93, and the variance ratio $\mathrm{Var}(\tilde{R}_o^g)/\mathrm{Var}(\Delta R_o) = 0.93$, indicating that SC debiasing retains 93% of the treatment variation.

Figure A.10: Heterogeneity in the Hours Effect (Women)

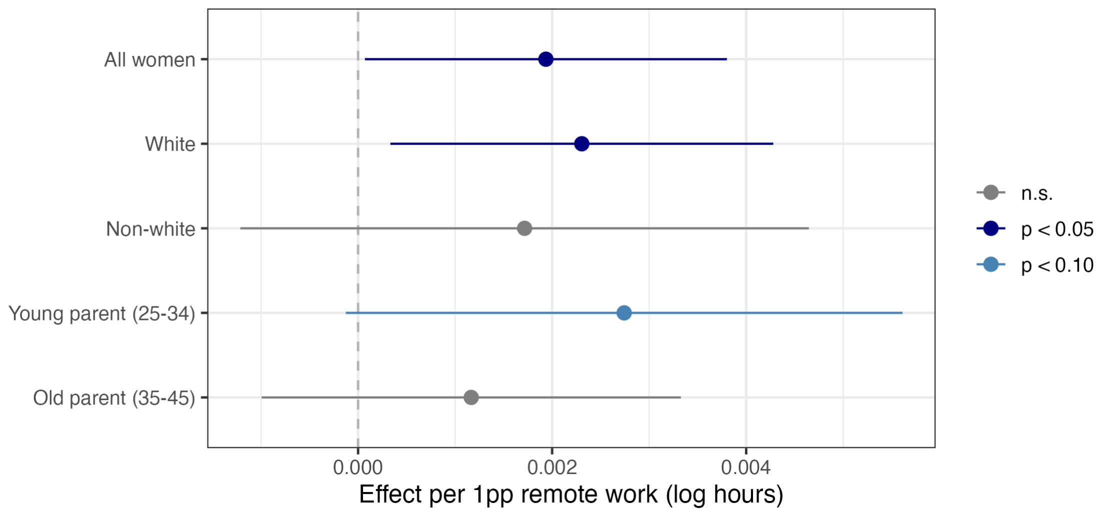

Notes: SC-DiD estimates of the effect of remote work on mothers' log hours, by subgroup. Continuous treatment ΔR_o, four child-age bins, separate matching, size-weighted. 95% CIs from randomization inference (1,000 permutations). Young parents are aged 25–34 at first birth; old parents are aged 35–45.

Figure A.11: Anticipatory Occupational Sorting into Remote Occupations

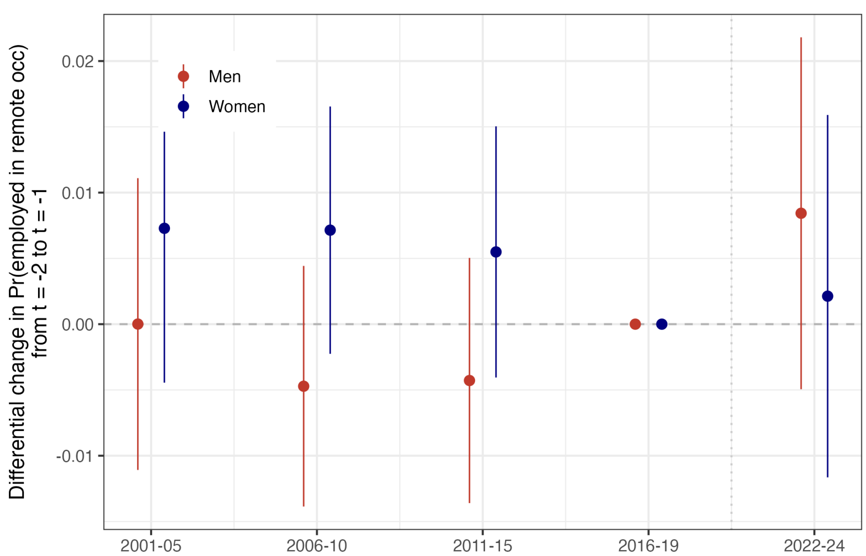

Notes: Each coefficient measures the differential change in Pr(employed in a remote occupation) from event time $t = -2$ to $t = -1$ for matched non-parents, relative to the reference period (2016–19). Non-parents at $t = -1$ are drawn from ACS year $T - 1$ and matched to new parents in year T on sex, education, marital status, race, birth cohort, and state; non-parents at $t = -2$ are drawn from ACS year $T - 2$ using the same matching cell. All ACS years 2020–2021 are excluded. Regression includes age, year, and education fixed effects, weighted by survey weights. 95% CIs from robust standard errors. If individuals anticipatorily sort into remote occupations before having children after COVID, the 2022–24 coefficient should be positive. For women, the post-COVID coefficient is 0.002 ($p = 0.762$); all coefficients are statistically insignificant for both genders.

Figure A.12: Child-Age-Specific Treatment Effects of Remote Work

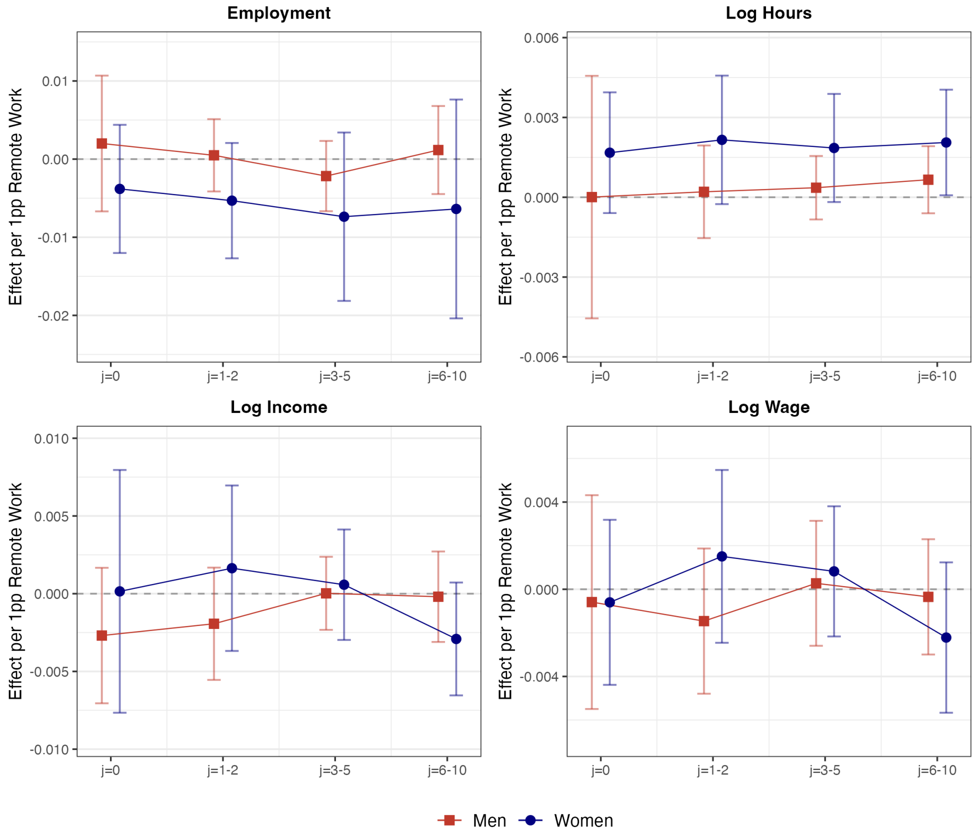

Notes: SC-DiD treatment effect of occupation-level remote work adoption (ΔR_o, continuous, per 1 pp) on the child penalty, estimated separately for four child-age bins: $j = 0$, 1–2, 3–5, and 6–10. Blue circles: women. Red squares: men. SC weights are the same as in Figure 3 (joint matching across employment and hours, size-weighted). 95% CIs from randomization inference (1,000 permutations). The hours effect for women is positive and similar in magnitude across all four child-age bins, indicating that the effect is not concentrated among mothers of very young children.

Figure A.13: Leave-One-Out Sensitivity of SC-DiD Estimates

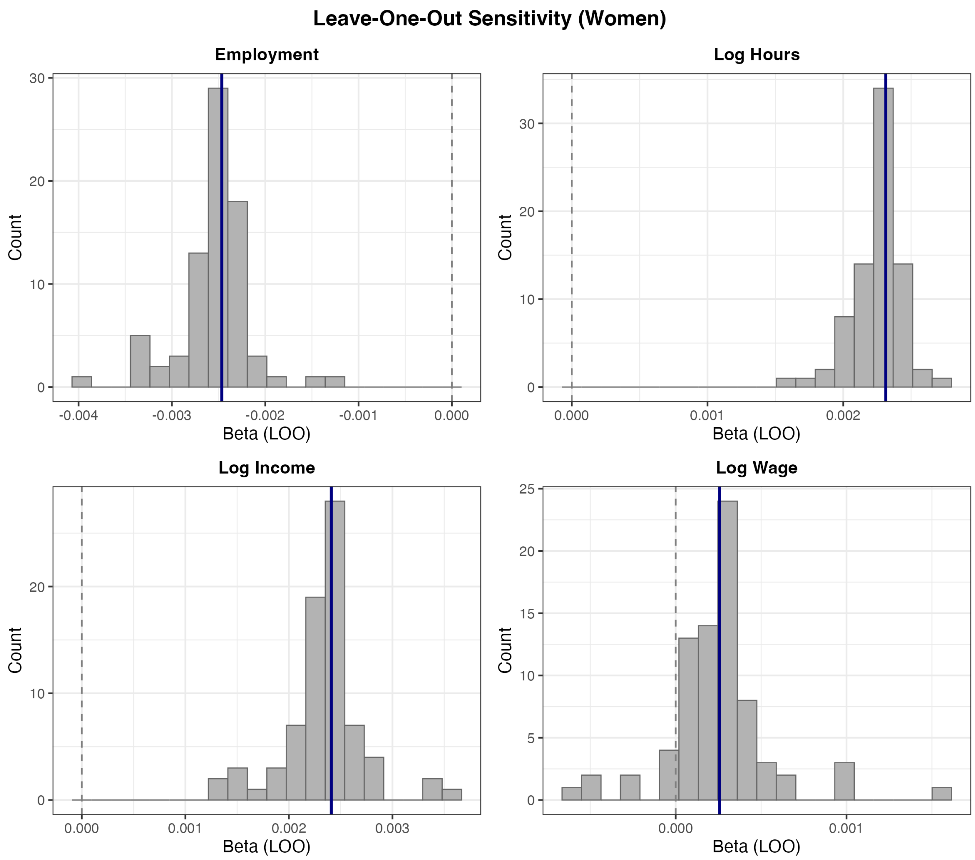

Notes: Each histogram shows the distribution of SC-DiD coefficients when one occupation is dropped from the sample and SC weights are re-solved for the remaining 76 occupations. The solid navy line marks the full-sample estimate (77 occupations). The dashed grey line marks zero. All estimates use the preferred specification: joint matching across employment and hours, 4 child-age bins, continuous treatment, size-weighted. For log hours, all 77 leave-one-out estimates are positive, ranging from +0.0015 to +0.0027 around the full-sample estimate of +0.0023. For log income, all 77 estimates are likewise positive, ranging from +0.0012 to +0.0036 around the full-sample estimate of +0.0024.

Figure A.14: CPS Telework Validation of Lightcast Treatment

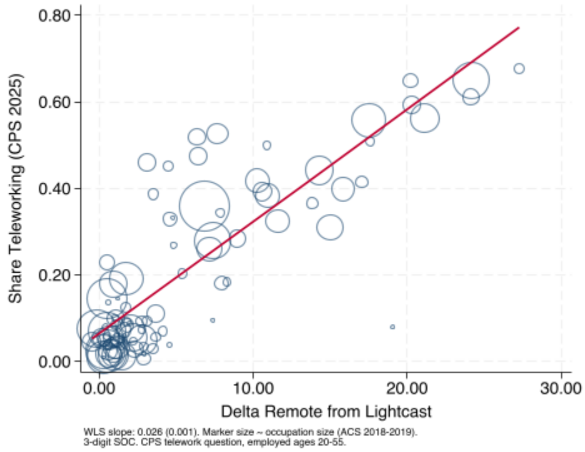

(a) Share teleworking, pooled

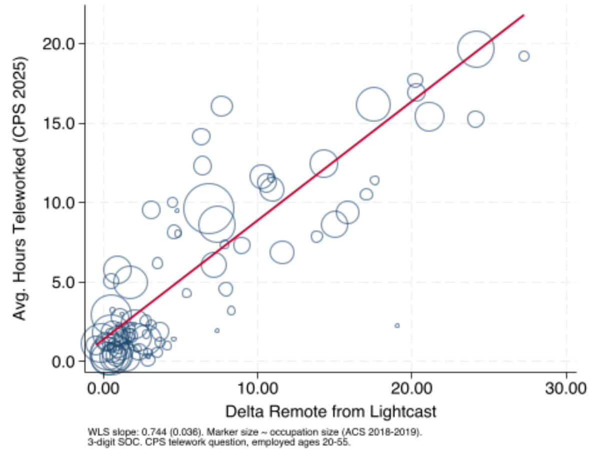

(b) Avg. hours teleworked, pooled

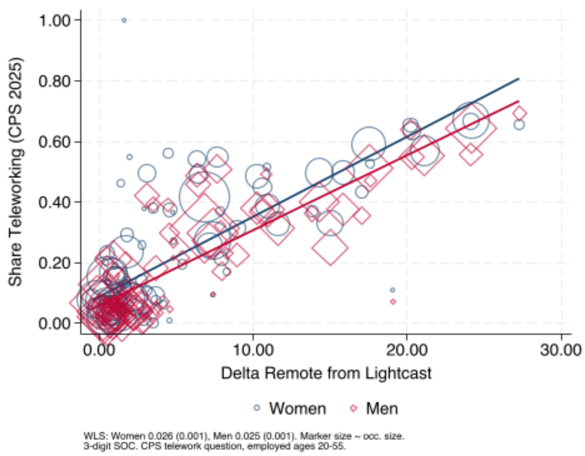

(c) Share teleworking, by gender

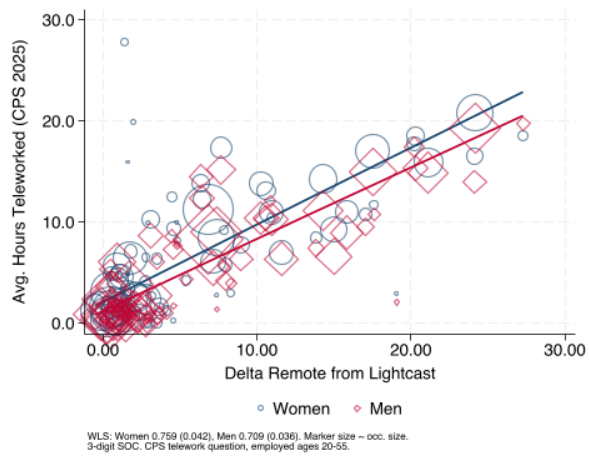

(d) Avg. hours teleworked, by gender

Notes: Each bubble is a 3-digit SOC occupation, sized by employment (ACS, 2018–2019). x-axis: ΔR_o, change in remote work share from Lightcast job postings (pp, 2017–19 vs. 2023–25). y-axis: realized telework intensity from the CPS monthly files, 2025, employed ages 20–55. Top row: all workers pooled. Bottom row: women (blue circles) and men (red diamonds) computed separately. WLS slopes reported in each panel. The strong positive relationship ($r = 0.82$ for the share, $r = 0.83$ for hours) confirms that employer-posted remote work in job ads translates into realized remote work arrangements.