# Industrial Automation, Earnings Parity, and Intra-Household Labor Supply

Claudio Costanzo*

September 2026

**Abstract**

Industrial robots erode men's wages more than women's, moving many couples closer to earnings parity. The added-worker effect predicts that when the main earner's wage falls, the other spouse works more. American time-use diaries show the opposite. Instrumenting local robot exposure with European adoption, I find that mothers in more exposed labor markets reduce market work by about two hours per week for each robot per thousand workers and shift that time to childcare and leisure, while fathers move the other way, more weakly. The withdrawal nearly triples where traditional gender norms are strong. To explain these patterns, I estimate a household model in which a breadwinner norm makes the wife's earnings worth less to the family as they approach her husband's, only 35 cents per dollar at parity. The model reproduces the reallocation, and the childcare margins identify the norm: as the wife's wage passes her husband's, her childcare turns up while his turns down, opposite reversals that bargaining power alone cannot generate. The estimated norm converts a shock that would narrow the gender gaps in work and childcare into one that widens both.

**Keywords:** Childcare; Gender identity penalties; Household labor supply; Industrial automation; Relative wages.

**JEL Codes:** J16, J22, J31, D13, O33.

*ECARES, Université Libre de Bruxelles. Contact: claudio.costanzo@ulb.be. I am grateful to Paula Gobbi, Bram De Rock, Fabio Blasutto, David Hémous, Luke Heath Milsom, as well as seminar participants at ULB–ECARES, KU-Leuven, UC3M, and UZH for their helpful comments.

# 1 Introduction

Industrial robots displace the routine, brawn-intensive tasks that men have traditionally performed, while raising the relative demand for the interpersonal tasks in which women hold the advantage (de Vries et al., 2020). The asymmetry is visible in the occupational distribution of exposure, where occupations with high robot exposure are predominantly male and those with little exposure predominantly female (Figure 1). The diffusion of robots has accordingly narrowed the gender wage gap by devaluing male-dominated routine labor (Ge and Zhou, 2020; Giuntella et al., 2025) and has lowered men's value in the marriage market (Anelli et al., 2024). This paper asks what that shift does inside households: how couples with young children reallocate market work, childcare, and leisure when automation erodes the husband's comparative advantage in the market.

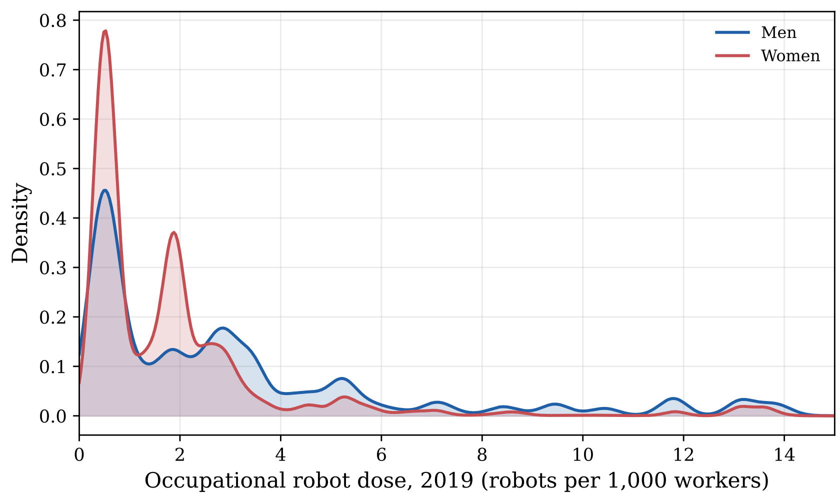


Figure 1: Kernel densities, by sex, of the occupational robot dose: each occupation's national 1990 industry mix over the nineteen IFR industries times 2019 U.S. penetration per 1,000 1990 workers (the measure of Online Appendix Table A35). Workers aged 18–64, 1990 five-percent Census sample (Ruggles et al., 2024), person-weighted; means 3.8 (men) and 2.3 (women) robots per thousand workers. Patent-based measures agree: Webb (2019) constructs occupational exposure from the textual overlap between O*NET task descriptions and robot-related patents and finds the same male tilt.

What a frictionless household would do with such a shift depends on the empirical shape of the shock. At Census scale, a unit of robot exposure lowers men's hourly wages by 0.36 percent and women's by 0.27 percent, and it does not reduce men's employment. The household's full income therefore falls while the wife's relative wage rises. For a shock of this shape, a frictionless model (Becker, 1965; Chiappori, 1992) makes two sign predictions and leaves one quantity open. Lower full income pushes both spouses toward more market work, the added-worker response (Stephens Jr., 2002). The wife's improved comparative advantage reallocates market work toward her and home production toward her husband. Her own hours are the quantity that remains open, because her own wage falls as well. Theory nonetheless pins down the direction of relative specialization. No combination of income, substitution, and bargaining effects moves market work toward the husband and childcare toward the wife when her relative wage has risen.

The data show precisely this excluded reallocation, and the paper establishes it in three steps. The first is cross-sectional. In American Time Use Survey (ATUS) diaries, couples' time allocation is sharply non-monotonic in the wife's share of potential wages, with a kink at earnings parity: her market work expands as her share approaches one half and contracts once it passes that threshold, her childcare traces the mirror-image V, and her husband's time use bends in the opposite direction.

The second step turns to causal evidence. Instrumenting local robot exposure with adoption in five European economies (Acemoglu and Restrepo, 2020), I find that greater exposure causes mothers to withdraw from market work into home production and leisure, by roughly 50 minutes per day per standard deviation of the shock, or about 16 minutes per robot per thousand workers, while fathers' diary minutes move in the opposite direction, more weakly and less robustly.

The third step identifies the channel. The female withdrawal is nearly three times larger in commuting zones with a high share of traditional religious adherents, and the geography of those zones is distinct from the geography of the shock. The withdrawal does not reflect displacement of the women themselves: it is absent precisely where robots reach women's own industries, occupations, and demographic cells, and it emerges instead as a response to the local male labor market.

A structural household model then quantifies the friction to which these facts point. Following the tradition of Akerlof and Kranton (2000) and Bertrand et al. (2015), a breadwinner norm enters as a wedge on the wife's effective earnings, and the wedge steepens as her share of potential wages approaches and passes parity. Spouses choose market hours cooperatively and childcare non-cooperatively, a timing that generates the corner solutions daily diaries display. I estimate the model by simulated method of moments on the levels and the kinked slopes of the four time-use profiles. The estimated model reproduces the reversals, and the childcare margins identify the wedge: as the wife's share passes parity her childcare turns up while her husband's turns down, and no wage-responsive bargaining weight reproduces that pair of opposite reversals, not even a fully cooperative one free to kink and jump at parity (Proposition 1). Because the wedge operates on the wife's return to market work, the husband adjusts only through household income and through the hours she relinquishes, and the model accordingly reproduces the weaker, mirror-image male response found in the causal estimates.

The friction the estimation recovers is large. At parity it removes roughly 65 percent of the consumable value of the wife's earnings, the scale required to overturn the added-worker response. With the wedge switched off, the same estimated household answers the same shock by having the higher-wage wife expand her market hours and out-work her husband. The norm therefore converts a shock that would narrow the gender gaps in work and childcare into one that widens both.

The same estimates price the policy response. Because the norm operates as an implicit tax on the wife's earnings, the tax code reaches it directly: a credit of ten cents on her dollar buys back a weekly hour of her market work for \$7 of fiscal outlay where the norm binds and for \$162 in the identity-neutral household, while the phase-out that family-based transfers impose destroys her hours at up to eighty times the frictionless excess burden. Wage insurance for displaced husbands protects female labor supply only through the norm's trigger, and public childcare raises female hours in both regimes, though the norm household routes half as much of the freed time into market work.

These results contribute first to the automation literature, which has measured the incidence of robots on wages, employment, and local labor markets (Acemoglu and Restrepo, 2020; Graetz and Michaels, 2018; Dauth et al., 2021). The estimates show that robots also reallocate time within households, and that the within-household incidence of the local shock is governed by norms.

They bear next on the evidence for gender identity and the aversion to a wife out-earning her husband (Brines, 1994; Bertrand et al., 2015). Relative to the earnings-share discontinuity, which later work has challenged on endogeneity and measurement grounds (Zinovyeva and Tverdostup, 2021; Hederos and Stenberg, 2022; Binder and Lam, 2022), the design brings plausibly exogenous variation in relative potential wages, together with a moderator whose geography is orthogonal to the shock's.

For structural household economics (Chiappori, 1992; Del Boca et al., 2014; Doepke and Kindermann, 2019), finally, the paper supplies a tractable formulation in which the norm enters the budget constraint instead of the utility function. Because the wedge bends marginal returns, it can be identified from the shape of time-use profiles alone, with the two parents' opposite childcare reversals as its signature.

Section 2 describes the data and documents the stylized facts. Section 3 constructs the exposure measure and the instrument, establishes the shock's effect on local wages, presents the causal estimates together with their robustness, and assembles the evidence on the mechanism. Section 4 develops and estimates the model, quantifies the distortion the norm imposes, and prices the policies that operate on it. Section 5 concludes.

# 2 Data and stylized facts

The analysis combines individual-level time diaries from the American Time Use Survey (ATUS) with data on the operational stock of industrial robots from the International Federation of Robotics (IFR). This section describes the two sources and then documents the cross-sectional facts that the causal design of Section 3 probes and the estimation of Section 4 targets.

**Time use.** The ATUS records how a nationally representative sample of Americans allocates the 24 hours of a specific day. Since 2003 the survey has been administered to a subset of households that have completed their final month in the Current Population Survey (CPS), so each diary is linked to a broad set of demographic and labor-market variables. The cross-sectional design prevents the tracking of individual transitions over time, but the diary captures daily allocations at a granularity that stylized labor-supply questions do not reach.

The primary sample is restricted to married or cohabiting households with at least one employed spouse and at least one child aged 10 or younger. This restriction follows Blundell et al. (2018) and concentrates the analysis on households in which the trade-off between market labor and intensive home production binds most tightly. Market work corresponds to the ATUS category "work and work-related activities"; childcare aggregates "Caring for and helping household children" with the activities related to household children's education and health. Because the local-labor-market design requires assigning each respondent to a commuting zone, the 2SLS sample is further limited to respondents whose county is identified in the public ATUS files. The Census Bureau suppresses county for respondents in smaller counties for confidentiality, so this restriction retains roughly 44% of ATUS respondents, those living in larger metropolitan commuting zones (about 190 of the 741 zones, of which 133–142 enter a given estimation column).

Table 1 reports descriptive statistics for this sample. The sample is highly educated, with tertiary-degree holders exceeding 44% among both men and women, yet its baseline time allocation follows a gendered division of labor: men average 262 minutes of daily market work against 132 for women, and women devote an average of 143 minutes to childcare, nearly double the 79 minutes recorded for men. The large standard deviations relative to the means are characteristic of time-diary data and reflect the variance between workdays and non-workdays captured in 24-hour snapshots. The discrepancy between usual weekly hours (44 for men, 23.5 for women) and diary minutes indicates that the sample includes a significant number of part-time working women and over-employed men. As the comparison with the broader population in Online Appendix Table A1 shows, the presence of young children compresses leisure for both spouses and redirects the household's time budget toward market work and active caregiving.

The outcomes of interest are the minutes a parent devotes to market work, childcare, and leisure, and only a time diary measures all three. Diaries are the reference instrument of the time-use literature: the respondent accounts for the 24 hours of a specific day, activity by activity, under an adding-up constraint, whereas stylized recall questions ("how many hours do you usually work?") elicit schedules and carry large and systematic errors (Juster and Stafford, 1991; Robinson and Bostrom, 1994; Bound et al., 2001; Kan and Pudney, 2008). Recall reports of hours cluster at round numbers and track the usual schedule, leaving little trace of the week just past (Robinson and Bostrom, 1994; Bound et al., 2001). Such reports register changes of job or contract, but not the day-level adjustments through which time actually worked changes at the margin: a day without market work, a shorter or longer working day, work outside the scheduled hours. Frazis and Stewart (2004) show that CPS recall hours and ATUS diary hours agree on average for the same reference week, with over- and under-statement offsetting across groups, so the diary adds no bias in levels while recording the realizations that recall questions smooth over; Section 3.3 shows that the responses documented here operate precisely on those realizations.

Table 1: Descriptive statistics: ATUS analysis sample (households with young children)

| | Females | | | Males | | |
|---|---|---|---|---|---|---|
| Variable | Obs | Mean | Std. Dev. | Obs | Mean | Std. Dev. |
| Leisure (min/day) | 24,748 | 226.80 | 164.97 | 22,021 | 262.00 | 195.74 |
| Childcare (min/day) | 24,748 | 142.94 | 147.38 | 22,021 | 79.19 | 113.27 |
| Market work (min/day) | 24,748 | 132.19 | 215.60 | 22,021 | 261.95 | 277.40 |
| Age | 24,748 | 35.49 | 6.78 | 22,021 | 38.05 | 7.54 |
| Hourly wage (USD) | 7,381 | 16.45 | 10.49 | 7,839 | 18.68 | 10.21 |
| Weekly earnings (USD) | 14,013 | 757.96 | 534.22 | 16,666 | 1,086.80 | 594.43 |
| Usual work hours per week | 23,048 | 23.56 | 19.61 | 20,463 | 44.00 | 14.24 |
| Less than secondary (%) | 24,748 | 7.85 | 26.90 | 22,021 | 8.72 | 28.22 |
| Secondary (%) | 24,748 | 19.12 | 39.33 | 22,021 | 22.17 | 41.54 |
| Some college (%) | 24,748 | 26.54 | 44.16 | 22,021 | 24.83 | 43.20 |
| Tertiary (%) | 24,748 | 46.48 | 49.88 | 22,021 | 44.28 | 49.67 |

*Notes:* The sample is restricted to married/cohabiting households with at least one employed spouse and at least one child aged $\leq$10. Leisure, childcare, and market work are ATUS diary minutes/day. Hourly wage, weekly earnings, and usual hours are derived from the linked CPS files. Education rows report the share (%) within each attainment category.

**Industrial robots.** Data on the operational stock of industrial robots come from the IFR, which defines an industrial robot as an "automatically controlled, reprogrammable, multipurpose manipulator, programmable in three or more axes." The definition identifies autonomous physical machines that perform motor tasks without a dedicated human operator and distinguishes them from traditional fixed-purpose automated machinery.

The gender asymmetry of the shock follows from the tasks these manipulators perform. Industrial robots substitute for brawn-intensive manual work such as welding, palletizing, and machining, concentrated in production occupations where men remain heavily overrepresented (Ge and Zhou, 2020; Cortes et al., 2024; Acemoglu and Restrepo, 2020). Where robots enter the more female-intensive corners of manufacturing, as in electronics assembly and inspection, they more often operate alongside workers in precision and oversight roles; and outside manufacturing, where the great majority of employed women work, adoption leaves the interpersonal and cognitive tasks of service work largely untouched. Robotization is therefore a gender-skewed shock to relative potential wages, devaluing the routine physical labor that men supply while leaving the tasks in which women are concentrated comparatively protected. The local wage estimates of Section 3.2 confirm the pattern, with men's wages falling by more than women's.

The U.S. operational stock grew slowly through the 1990s and accelerated sharply after 2000 (Online Appendix Figure A1), so the 2003–2019 diary window covers the steep part of the diffusion. Two measurement issues arise for the United States. Sectoral disaggregation is consistently available only from 2004 onward, because earlier years contain a high share of robots not attributed to a specific industry; the analysis uses detailed data for 13 manufacturing sub-sectors and six non-manufacturing categories, and allocates the robots that the IFR classifies as "unspecified" proportionally to the distribution of the classified stock within the country, following Acemoglu and Restrepo (2020).[1] And because the IFR reports stocks only at the national level, local exposure is constructed by mapping national industry trends to the 1990 industrial composition of U.S. commuting zones, as Section 3 details; the mapping converts national technological change into local labor-market variation driven by historical specialization in robot-intensive industries.

## 2.1 Stylized facts on wage shares and time use

In the cross-section, couples' time allocation is sharply non-monotonic in the wife's share of household wages: as her share approaches and passes her husband's, both spouses reallocate

[1]The categories are Agriculture and Fishery, Automotive, Construction, Electronics, Food and Beverages, Furniture, Basic Metals, Machinery, Metal Products, Non-metallic Minerals, Mining, Paper, Petrochemicals, Research, Services, Textiles, Utilities, and Non-automotive Vehicles.

market work, childcare, and leisure in distinct, kinked patterns. This subsection documents these facts, which the estimation of Section 4 later targets; Online Appendix Section B shows that family-oriented beliefs in the International Social Survey Programme (ISSP) bend at the same threshold.

The sample consists of married or cohabiting couples under age 64 with at least one child aged ten or younger, with at least one spouse employed and active job seekers excluded, on non-search diary days, the same configuration as the structural estimation sample of Section 4. Following Bertrand et al. (2015), hourly wages are computed by dividing weekly earnings by weekly hours worked.[2] For non-workers, a potential wage is imputed as the mean wage of employed individuals sharing the same education level, five-year age bracket, state of residence, and survey year. Including non-workers keeps both margins of labor supply in view, hours and participation.

The wife's relative position is measured by her share of hourly wages, $s = w_f/(w_f + w_m)$. Earnings shares would be the natural alternative, and they fail for two reasons. Because the outcome variables include market hours, realized earnings are mechanically endogenous: a woman who reduces her working time to comply with the breadwinner norm lowers her own share, sorts into lower bins, and obscures the behavioral threshold. Potential earnings, in turn, embed the baseline gender gap in weekly hours: reaching potential-earnings parity requires the wife's hourly wage to be nearly double her husband's, so an unconditional cross-section in that share is selected. The hourly wage-rate share avoids both distortions. Online Appendix Figure A2 plots realized weekly earnings against this share: the wife's share of household earnings rises from 0.25 below parity to 0.37 above it, the husband's earnings fall across the threshold while hers rise, and household earnings decline only mildly. The share itself has no mass point at 0.5 beyond the 6.6 percent of couples with equal imputed potential wages, and excluding those couples leaves the profiles unchanged (Online Appendix Table A3).

The resulting profiles are the paper's central stylized facts.

**Time-use patterns.** *As the wife's share of the couple's potential hourly wages approaches and passes one half, her market work rises and then falls while her childcare falls and then rises; her husband's childcare and leisure rise and then fall, and his market work rises once her share passes one half.*

Figure 2 displays the patterns: average daily minutes of market work, childcare, and leisure against the wife's wage share, with separate linear trends fitted on either side of the parity threshold. For women (left column), the narrow confidence bands around the work and childcare fits show that the slope reversals at 0.5 are systematic. Her working time increases

[2] Observations implying hourly wages above \$100 are dropped due to top-coding truncation.

as her wage share rises toward parity and reverses past the threshold, an inverse-V, and her childcare traces the opposite, V-shaped path. Female leisure is flat. The breadwinner norm pushes the wife's time toward home production, while her rising wage share strengthens her bargaining position and lets her claim more private leisure (Chiappori, 1992; Browning and Chiappori, 1998); in the cross-section the two channels roughly offset, and her leisure barely moves.

A corresponding reversal characterizes men's time use (right column). Male market work is flat or declines slightly as the wife contributes more, and rises steeply once she out-earns him; male childcare and leisure follow an inverse-V, increasing up to the parity point and declining beyond it, with the decline the sharper arm. The wider confidence intervals for men reflect greater idiosyncratic variation in how husbands adjust as the couple approaches and crosses the equal-earnings boundary. The patterns suggest that once the traditional breadwinner hierarchy is threatened, households deviate from standard specialization, with men increasing market effort and women taking on more non-market tasks to compensate for the norm violation. The location of the kink is an empirical object in its own right: re-estimating the two-segment profiles with a freely located knot places the reversal at parity where the data can identify it, most sharply for maternal childcare ($\hat{k} = 0.52$, bootstrap 95% confidence interval $[0.44, 0.58]$), and imposing the kink at 0.5 costs less than 0.04% of fit in every profile (Online Appendix Figure A6; Section 4).

These profiles bear on a contested literature. Bertrand et al. (2015) documented a discontinuity in the distribution of the wife's share of couple earnings at one half and read it as aversion to the wife out-earning the husband. Subsequent work has complicated that reading. The discontinuity is produced largely by a point mass of equal-earning couples, which Zinovyeva and Tverdostup (2021) trace to coworking spouses and Hederos and Stenberg (2022) find in Swedish registers as well; Binder and Lam (2022) show that matching models without any norm reproduce the distribution; and reported earnings are manipulated around the threshold (Heggeness and Murray-Close, 2019). The evidence here differs in three respects. The running variable is the ratio of hourly wages, with potential wages imputed for non-workers, so realized earnings never enter it: there is no mechanical mass at parity, and a change in the wife's hours does not move the share by construction. The object is a reversal of slopes on either side of parity, a different shape from the contested jump in the density, and the free-knot estimates of Online Appendix Section G place the kink at 0.52 with a confidence interval of $[0.44, 0.58]$. And the profiles are unchanged when couples within one percent of equal wages are excluded (Online Appendix Table A3). Section 4 places the estimated friction next to the policy-experiment evidence on norms (Lippmann et al., 2020; Ichino et al., 2026; Grönqvist et al., 2026).

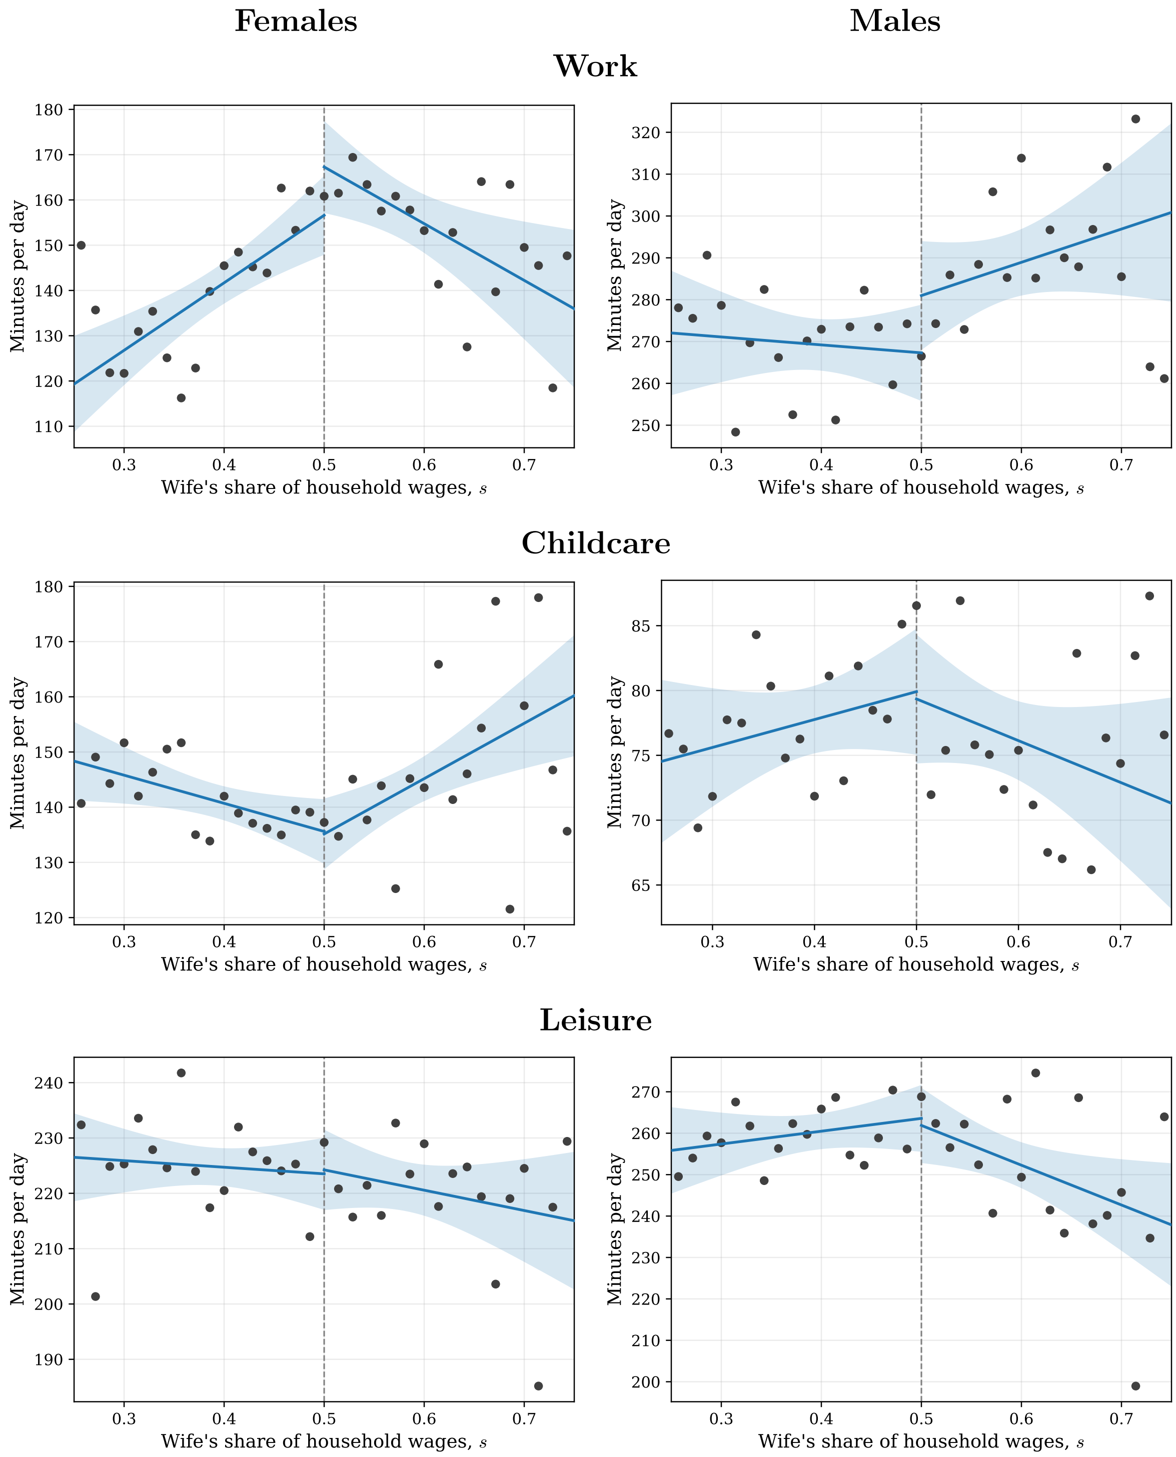


Figure 2: Relationship between the wife's wage share $s = w_f/(w_f + w_m)$ (horizontal axis) and average daily minutes devoted to work, childcare, and leisure (vertical axis), for women (left column) and men (right column). Markers are local averages within wage-share bins; solid lines are linear fits allowed to kink at parity ($s = 0.5$, dashed vertical line); shaded bands are 95% confidence intervals. The sample and the wife's-wage-share construction are the same as those underlying the structural estimation targets of Section 4 (26,654 respondents over the interior support $0.25 < s < 0.75$), so the descriptive figure and the estimated moments share one sample.

These descriptive patterns cannot, on their own, establish behavior. Spouses adjust their labor supply to unobserved preferences and local conditions, and women with strong career attachment sort into particular marriages. Section 3 therefore turns to industrial automation as an exogenous shock to relative wages. Because robots devalue male-dominated routine work in the lower-middle of the wage distribution, the shock lands on households in which the wife's relative wage is already high and pushes them further into the region where the profiles bend. If the kink comes from behavior instead of sorting, exposed couples should reallocate the way the cross-section suggests; the next section shows that they do.

# 3 The effect of robots on time use

This section estimates the causal effect of robot exposure on couples' time use. It constructs the local exposure measure and its European instrument, verifies that the shock reaches local wages the way the mechanism requires, presents the time-use estimates together with the checks they survive, and closes with the evidence on the mechanism behind them.

## 3.1 Empirical design

Local exposure to robots follows a Bartik-style construction (Acemoglu and Restrepo, 2020):

$$\text{Robots}_{ct} = \sum_j \frac{\text{Emp}_{jc}^{1990}}{\text{Emp}_c^{1990}} \times \frac{\text{StockRobots}_{jt}}{\text{Emp}_j^{1990}}, \tag{1}$$

where $\text{Emp}_{jc}^{1990}/\text{Emp}_c^{1990}$ is the 1990 industrial specialization of commuting zone $c$, so that regional exposure is anchored to employment shares predating the surge in automation. Throughout, $j$ indexes industries, $e$ indexes the EU5 economies, and $st(c)$ denotes the state of commuting zone $c$; the letter $s$ remains reserved for the wife's wage share.

Even with historical shares, U.S. industry-level adoption ($\text{StockRobots}_{jt}$) might be correlated with local shocks. Following Acemoglu and Restrepo (2020), I therefore instrument U.S. exposure with adoption in five advanced European economies (Denmark, Finland, France, Italy and Sweden, the EU5), the countries for which the IFR reports industry-level stocks from 1993; the same strategy is used by Graetz and Michaels (2018) and Dauth et al. (2021). The instrument is

$$\text{Robots}_{ct}^{IV} = \frac{1}{5} \sum_{e \in \text{EU5}} \left( \sum_j \frac{\text{Emp}_{jc}^{1970}}{\text{Emp}_c^{1970}} \times \frac{\text{StockRobots}_{jt}^{e}}{\text{Emp}_j^{e,1990}} \right), \tag{2}$$

where the 1970 employment shares further decouple the instrument from recent U.S. economic shifts. The assumption borrowed with the strategy is that, conditional on the fixed effects, EU5 adoption in an industry is unrelated to the household labor-supply shocks of the U.S. commuting zones specialized in that industry. Three observations support it. First, the EU5 series lead the U.S. series in time: the cross-industry pattern of EU5 adoption over 1995–2004, before the IFR reports any U.S. sectoral data, predicts the pattern of U.S. adoption over 2004–2019 (a correlation of 0.85 across the nineteen industries, 0.45 in ranks), with automotive first in both. Shifters measured three and five years before the diary year, which cannot respond to contemporaneous U.S. shocks, reproduce the baseline estimates (Online Appendix Table A5). Second, coordinated adoption by multinationals operating on both sides of the Atlantic strengthens the first stage without threatening exclusion: the state-by-year effects absorb any aggregate component of adoption and of U.S. time use, so the exclusion restriction fails only if European adoption in an industry responds to the household labor-supply shocks of the particular U.S. zones specialized in it. Third, the five national series are not interchangeable, and their differences are informative. French and Italian automotive penetration declined over 2004–2019 while U.S. penetration rose, so within the fixed effects these two series enter the first stage with a negative sign, and the relevance of the average comes from the Nordic series. Used one at a time, all five nonetheless deliver the female withdrawal and the leisure responses, with reduced forms that change sign together with the first stage, as they must if the instrument operates through U.S. exposure (Online Appendix Tables A5, A6 and A7).[3]

The estimating equation relates daily time use to instrumented exposure via two-stage least squares (2SLS):

$$Y_{ict} = \alpha + \beta\,\widehat{\text{Robots}}_{ct} + \eta' C_{ict} + \gamma_c + \delta_{st(c)\times t} + \varepsilon_{ict}, \tag{3}$$

where $Y_{ict}$ is the daily minutes spent in market work, childcare, or leisure by individual $i$ in commuting zone $c$ and year $t$, and the first stage is

$$\text{Robots}_{ct} = \pi\,\text{Robots}^{IV}_{ct} + \eta' C_{ict} + \gamma_c + \delta_{st(c)\times t} + u_{ct}. \tag{4}$$

The vector $C_{ict}$ contains temporal controls (day of week, month, holiday), individual demographics (age, education), and household characteristics (number and age of children, spouse's age and education); commuting-zone fixed effects ($\gamma_c$) absorb persistent regional differences and state-by-year fixed effects ($\delta_{st(c)\times t}$) absorb broader economic trends. The coefficient $\beta$

[3] The IFR counts industrial robots as defined by ISO 8373 (automatically controlled, reprogrammable, multipurpose manipulators); the humanoid and mobile platforms piloted from 2025 fall outside both the definition and the sample period.

captures the average change in time use for a unit increase in local robot exposure, roughly one additional robot per thousand workers. All regressions are estimated with ATUS survey weights, and standard errors are clustered at the commuting-zone–year level; to guard against cross-regional sectoral shocks, Section 3.4 adds exposure-robust standard errors (Adão et al., 2019). Active job seekers are excluded throughout, which isolates the structural reallocation of time from short-term unemployment effects.

The identifying variation is highly concentrated: the automotive industry carries roughly 56% of the Rotemberg weight, and the effective number of independent industry shocks is only about three (Section 3.4, Online Appendix Table A11). The design is therefore best understood as a *shares*-based one in the sense of Goldsmith-Pinkham et al. (2020); it does not lean on a large number of quasi-random shocks (Borusyak et al., 2022, 2025; Adão et al., 2019). The operative exogeneity condition is accordingly on the historical shares: validity requires that 1970 and 1990 industrial specialization, above all the automotive share, be uncorrelated with unobserved regional trends in gender norms or female labor supply, conditional on the controls and fixed effects. Table 2 assesses the condition directly through a share-balance panel and a pre-trend placebo; a leave-automotive-out instrument, specifications that let 1990 characteristics of the commuting zone follow their own time paths, and the full Rotemberg and AKM diagnostics follow in Section 3.4. The state-by-year effects already absorb every state-level trend, union density, policy, and state-level norms among them, so only within-state differences across zones identify $\beta$. I read the estimates as the response of households to an exogenous, automation-driven erosion of the husband's comparative advantage.

The resulting exposure averages 3.14 robots per 1,000 workers in the sample and varies by an order of magnitude across commuting zones, from service hubs such as New York City (1.29) to manufacturing centers such as Grand Rapids and Detroit (approaching 6); the automotive sector's 1990 employment share is roughly fourteen times higher in the most exposed areas than in the least (Online Appendix Table A2). Figure 3, panel (a), maps the geography of this treatment. Exposure concentrates in the Manufacturing Belt, above all the automotive corridor of Michigan, northern Ohio and Indiana, and Wisconsin, with secondary clusters in the upper South's newer auto plants, while the coasts and the service-oriented Sun Belt are lightly exposed. This geography matters for identification twice over: it shows the variation the shares-based design exploits, and, set against the map of traditional religious adherence beside it (panel (b)), it displays the orthogonality that Panel A of Table 2 quantifies, since the automotive belt and the culturally conservative belt are different places.

Table 2: Identification checks: share balance and pre-trends

| **A. Correlation of the identifying shares with 1990 commuting-zone characteristics** | | |
|---|---|---|
| | 1970 automotive share | EU5 instrument |
| Female share | +0.144 (0.037) | +0.213 (0.047) |
| College share | −0.159 (0.059) | −0.178 (0.062) |
| High-school share | +0.141 (0.056) | +0.218 (0.060) |
| Manufacturing share | +0.329 (0.055) | +0.484 (0.052) |
| Log population | +0.161 (0.076) | +0.202 (0.067) |
| Black share | +0.011 (0.059) | −0.003 (0.061) |
| Traditional religious adherence | −0.175 (0.030) | −0.204 (0.037) |
| **B. 1990–2000 placebo: change in the outcome on the later change in exposure** | | |
| | Coefficient | $t$-statistic |
| Gender employment gap | +0.001 | 0.4 |
| Gender wage gap | +0.001 | 0.5 |
| Female labor-force participation | +0.012 | 5.5 |
| Male labor-force participation | +0.013 | 5.2 |
| Mean wife's wage share (couples) | +0.001 | 1.7 |
| Share of couples within 0.1 of parity | +0.001 | 0.2 |

*Notes:* Panel A: cross-commuting-zone correlations of the high-weight 1970 automotive employment share and of the EU5 instrument with predetermined 1990 characteristics, with bootstrap standard errors over commuting zones in parentheses (2,000 draws; Online Appendix Table A12). The shares load on manufacturing intensity, as any Bartik measure must; the traditional-adherence correlation is significantly *negative*. Panel B: the 1990–2000 change in each outcome, from the 1990 and 2000 decennial Census microdata collapsed to commuting zones, regressed on the 2004–2019 change in robot exposure with the 1990 covariates and Census-division fixed effects of the long-difference design; the gap and share placebos are null, while participation rose for both sexes by the same amount in soon-to-be-exposed zones, a common trend that differences out of every gendered contrast and runs against a female withdrawal.

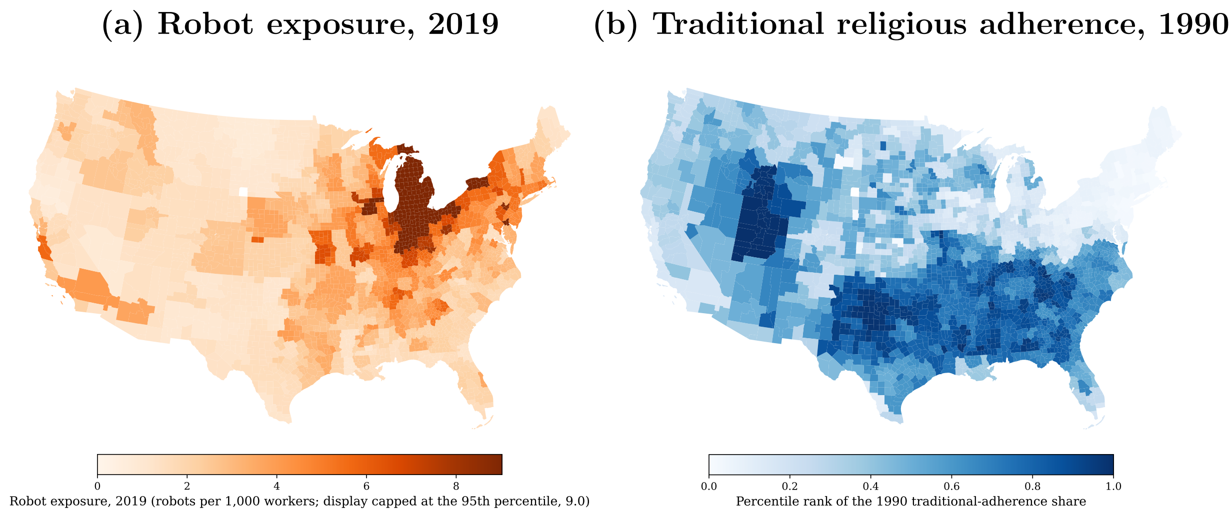


Figure 3: Commuting-zone robot exposure and traditional religious adherence. Panel (a): commuting-zone robot exposure in 2019 (robots per 1,000 workers), the terminal year of the panel; counties are colored by their commuting zone's exposure, with the display scale capped at the 95th percentile (9.0). Exposure concentrates in the Manufacturing Belt, led by the automotive corridor around Michigan. Panel (b): the 1990 share of traditional religious adherents. The two maps are distinct: the automotive belt is *less* traditional than the rest of the country (correlation $-0.18$), which lets Section 3.5 interact the shock with local norms without conflating treatment and moderator.

## 3.2 The shock in the labor market

Before turning to time use, I verify that the shock reaches wages the way the mechanism requires. In stacked commuting-zone long differences of the ATUS sample itself, a unit of exposure lowers men's hourly wages by 6.1 percent over three years and narrows the gender wage gap, with the effect attenuating over five and seven years (Online Appendix Table A14). The American Community Survey, with eighteen million working-age adults, confirms the sign and adds the employment margin (Table 3). Over 2005–2017, exposure lowers men's wages by 0.36 percent per robot and women's by 0.27, so the wife's relative wage rises, while men's employment does not fall. The employment coefficient cannot separate labor-market attachment from compositional change, and it is not used to identify a male-effort response, which rests on the time-use evidence.

The channel does not operate couple by couple: in 2.3 million ACS couples, exposure raises the share of wives out-earning their husbands by 0.13 percentage points per robot, so that a three-robot shock moves well under one percent of couples across parity (Online Appendix Table A16). The time-use response is therefore read as households reallocating when the local labor market moves against the male-breadwinner occupation mix, with an intensity governed by the local salience of the norm (Table 7) and by the couple's position in the wage-share distribution (Section 2.1). Accordingly, the structural model is identified off the cross-sectional time-use profile in $s$; a first stage of the shock on $s$ plays no role in it.

Table 3: Census-scale first stage: robot exposure on wages and employment by sex (ACS, 2SLS long difference)

| | No controls | | + 1990 controls, division FE | |
|---|---|---|---|---|
| Outcome (long diff 2005–07 → 2015–17) | Coef. | (SE) | Coef. | (SE) |
| Male ln hourly wage | −0.0046*** | (0.0007) | −0.0036*** | (0.0012) |
| Female ln hourly wage | −0.0026*** | (0.0009) | −0.0027*** | (0.0010) |
| Wage gap (male − female) | −0.0019*** | (0.0006) | −0.0010 | (0.0006) |
| Male employment rate | +0.0018*** | (0.0004) | +0.0020*** | (0.0004) |
| Female employment rate | +0.0001 | (0.0004) | −0.0005 | (0.0004) |
| Employment gap (male − female) | +0.0017*** | (0.0004) | +0.0024*** | (0.0005) |

*Notes:* 2SLS long differences across all 722 commuting zones from pooled ACS 2005–2007 and 2015–2017 (Ruggles et al., 2024), roughly 18.3 million working-age adults (24–64). Each commuting zone's change in the sex-specific mean log hourly wage / employment rate is regressed on the change in robot exposure, instrumented by the EU5 shifter; ACS persons are mapped to commuting zones with population-weighted PUMA crosswalks. Right panel adds 1990 commuting-zone covariates and Census-division fixed effects, as in Online Appendix Table A14. Weighted by base-period population; heteroskedasticity-robust standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

## 3.3 Results

Table 4 reports the response of daily time use to a one-unit increase in robot exposure, roughly one additional robot per 1,000 workers, across three subsamples that separate the margins of labor supply. Columns 1 and 4 condition on an employed spouse, so the respondent's own time use may adjust along the extensive margin, including non-employment. Columns 2 and 5 condition on an employed respondent, allowing extensive-margin adjustments by the spouse. Columns 3 and 6 restrict the sample to dual-earner households, isolating the intensive margin for both partners. In every column, $\text{Robots}_{ct}$ is instrumented with the EU5 shifter.[4]

For men, diary minutes of market work rise: by approximately 13 minutes when non-employment is allowed (column 1), by roughly 9 minutes conditioning on the respondent's employment (column 2), and by 14 minutes in two-earner households (column 3), with leisure falling by roughly 9 minutes in the latter group. The male estimates are about two-thirds the size of the female ones in absolute value and, as Section 3.4 shows, less robustly identified; I read them as the weaker mirror image of the female response and give them no headline role.

For women, the signs reverse and the estimates are robust across subsamples: market work falls by approximately 16 minutes per day in the full sample (column 4), 19 minutes

[4] In all 2SLS specifications, the Kleibergen–Paap $F$-statistic (Kleibergen and Paap, 2006) for the first stage exceeds the critical values for weak identification: per column, $F$ ranges from 334 to 409 (reported in Table 4). The first-stage estimates themselves are in Online Appendix Table A4.

conditioning on the respondent's employment (column 5), and 16 minutes among dual earners (column 6). That the withdrawal is, if anything, larger among women who remain employed indicates that the response is not only an exit from the labor force but also a scaling back of hours within employment, cutting overtime or moving toward part-time schedules, as the household reallocates when the husband's labor-market position weakens. The diary pins down the margins: the probability of any market work on the diary day falls by 2.1 to 2.7 percentage points per robot (3 to 4 on weekdays), and the working day of those who work shortens by 15 to 17 minutes on weekdays, with men the mirror image at the ten-percent level (Online Appendix Table A17). These are realized days, the margin that a usual-hours question smooths over and the diary records (Section 2).

The withdrawn market time goes to the home. Women's childcare rises by 7 to 10 minutes across columns 4–6 and their leisure by 6 to 11 minutes, while men's leisure falls by 9 minutes among dual earners as their market minutes rise. The opposite leisure movements fit the collective reading: as the wife's relative potential wage rises it also acts as a distribution factor (Chiappori, 1992; Blundell et al., 2005), letting her capture more of the household's non-market time even as she absorbs more of its childcare. In magnitudes, moving from a bottom-decile zone such as New York City to the top-decile automotive heartland implies a female withdrawal on the order of an hour and a half per day; because the specification is linear in exposure, such figures are best read within the interquartile range of the shock. They are also large relative to wage-elasticity benchmarks: no own- or cross-wage substitution elasticity in the conventional 0.2–0.6 range delivers responses of this size to wage movements of a few percent (Chetty et al., 2011). The friction estimated in Section 4 is sized to fill that gap, and the usual-hours estimates of Section 3.4, at 1.0–1.2 hours per week, corroborate the diary scale.

Overall, robot exposure drives a gendered reallocation of time within couples. The standard income effect would have women increase their labor supply to offset the male wage loss; instead, women shift out of market work into childcare and leisure, and men's diary minutes move the other way. This joint pattern, market work moving toward the husband and childcare toward the wife as her relative wage rises, is the one a frictionless household cannot produce; whether a bargaining weight that rises with the wife's relative wage could produce it on its own is tested directly in Section 4.4. Because the effects are heavily pronounced even in dual-earner households, they point to a renegotiation of hours within the marriage; mechanical displacement alone would not concentrate there. The ATUS collects one diary per household, so the columns identify gender-specific average responses among partnered respondents, and the within-couple reallocation is inferred from the gendered pattern of the estimates. The one within-household margin the data carry, the spouse's employment and usual weekly hours as reported by the respondent, does not move (Online Appendix Table A23). The structural model of Section 4 supplies the joint accounting a single diary cannot, validated against the spouse's usual hours across the couple's wage share (Section 4.3).

## 3.4 Robustness

Two questions remain once the design and the first stage are in place: whether the time-use estimates survive changes of specification and of inference, and whether the response is a by-product of measurement or of concurrent shocks. This subsection takes both in turn; full outcome sets and further detail are in Online Appendix Section D.

**Specification and inference.** Table 5 collects the specification and inference checks for the headline market-work estimates, one table row per check, with the full outcome sets in the Online Appendix. The female withdrawal is not a weekend-composition artifact: on weekday diaries it grows to 26–29 minutes, as a labor-supply response should when working days are at stake. It survives instrumenting with the EU5 shifter measured three years before the diary. Reweighting for shock-induced selection into the couple sample leaves the estimates essentially at the baseline; the reweighting matters because automation moves marriage and fertility (Anelli et al., 2024), so the sample itself is potentially treated. One row probes the shares-based identification directly. Letting six predetermined 1990 characteristics of the commuting zone follow their own year-specific paths leaves the female withdrawal intact and larger, but absorbs the male increase, to between $-2$ and $+2$ minutes with the first stage intact. No single characteristic's trend does this, and the estimate stays at 10–13 minutes with any one of them, but the two manufacturing-share trends together absorb it; the male intensive-margin response is therefore identified less securely than the female one. The inference rows keep the baseline point estimates and change the standard errors. Clustering by commuting zone moves only men's dual-earner work, to $p = 0.053$. Exposure-robust (AKM) inference with $t$ critical values at $g_{\text{eff}} - 1$ degrees of freedom (Adão et al., 2019) is the conservative correction for a design with about three effective industry shocks: it preserves the female withdrawal in the broad and employed samples but not in the smallest dual-earner cell, puts the male increase at the five-to-ten percent level, and does not preserve women's childcare increase, the one margin identified only under conventional clustering.

Table 4: Effect of robots on daily minutes of work, childcare, and leisure (2SLS estimates).

| | Males | | | Females | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| **Daily minutes work** | | | | | | |
| Robots | 12.734** | 8.967** | 13.953** | -15.949*** | -18.510*** | -16.361*** |
| | (5.834) | (3.941) | (6.306) | (4.767) | (5.749) | (6.264) |
| Mean dep. var. | 260.2 | 281.7 | 278.6 | 134.3 | 218.8 | 216.5 |
| Observations | 4,100 | 6,593 | 3,812 | 7,194 | 4,891 | 4,400 |
| R-squared | 0.426 | 0.443 | 0.476 | 0.286 | 0.422 | 0.433 |
| **Daily minutes childcare** | | | | | | |
| Robots | 0.854 | 2.466 | 1.771 | 6.579** | 9.686*** | 10.092*** |
| | (2.260) | (2.031) | (2.214) | (2.627) | (2.815) | (3.030) |
| Mean dep. var. | 85.6 | 80.8 | 82.5 | 148.0 | 118.6 | 120.6 |
| Observations | 4,100 | 6,593 | 3,812 | 7,194 | 4,891 | 4,400 |
| R-squared | 0.229 | 0.180 | 0.253 | 0.239 | 0.231 | 0.246 |
| **Daily minutes leisure** | | | | | | |
| Robots | -5.494 | -3.772 | -8.935*** | 6.205** | 10.996*** | 9.746*** |
| | (3.930) | (3.184) | (3.230) | (3.057) | (3.062) | (3.659) |
| Mean dep. var. | 260.7 | 251.8 | 252.6 | 221.3 | 204.5 | 204.7 |
| Observations | 4,100 | 6,593 | 3,812 | 7,194 | 4,891 | 4,400 |
| R-squared | 0.289 | 0.259 | 0.313 | 0.208 | 0.305 | 0.314 |
| KP first-stage $F$ | 353 | 380 | 368 | 409 | 408 | 334 |
| Commuting zones | 136 | 139 | 135 | 142 | 134 | 133 |
| Spouse employed | ✓ | | ✓ | ✓ | | ✓ |
| Respondent employed | | ✓ | ✓ | | ✓ | ✓ |

*Notes:* Sample includes married or cohabiting couples with at least one spouse working and children aged 10 or less. Columns 1 and 4 capture the extensive margin of the respondent (spouse is employed). Columns 2 and 5 capture the extensive margin of the spouse (respondent is employed). Columns 3 and 6 restrict to dual-earner households, capturing only the intensive margin of labor supply. Control variables in all specifications include commuting-zone and state-by-year fixed effects, spouses' age and education, number of children (total and under age 5), ages of the youngest and oldest child, day-of-week and month-of-survey fixed effects, and a holiday indicator. Family income is excluded as a post-treatment "bad control," since the shock operates partly through household earnings. Mean dep. var. is the unweighted mean of the outcome in each column's estimation sample. The Kleibergen–Paap row reports the just-identified first-stage $F$ on the EU5 shifter; it is common to the three outcome panels because they share each column's estimation sample, and the first-stage estimates themselves are in Online Appendix Table A4. Standard errors clustered at the commuting-zone–year level are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: Robustness of the market-work estimates: specification and inference (2SLS)

| | Males | | | Females | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Baseline | 12.734** | 8.967** | 13.953** | -15.949*** | -18.510*** | -16.361*** |
| | (5.834) | (3.941) | (6.306) | (4.767) | (5.749) | (6.264) |
| Weekday diaries only | 17.904** | 13.051** | 18.263** | -26.549*** | -28.755*** | -26.396*** |
| | (7.574) | (5.915) | (8.206) | (6.579) | (7.202) | (7.897) |
| 1990 characteristics × year FE | 1.144 | -1.640 | -0.646 | -19.841*** | -18.326*** | -17.808** |
| | (6.549) | (4.679) | (6.953) | (5.488) | (7.001) | (7.419) |
| EU5 shifter lagged 3 years | 17.686*** | 11.168*** | 18.973*** | -15.906*** | -15.026** | -15.793** |
| | (4.615) | (3.500) | (4.899) | (4.495) | (6.183) | (6.678) |
| Couple-sample selection (IPW) | 12.514** | 5.782 | 12.973** | -16.053*** | -19.490*** | -17.808*** |
| | (5.617) | (4.368) | (6.335) | (4.998) | (5.718) | (6.286) |
| *Baseline point estimates under alternative inference* | | | | | | |
| Commuting-zone clusters | 12.734** | 8.967** | 13.953* | -15.949*** | -18.510*** | -16.361*** |
| | (6.156) | (4.294) | (7.159) | (4.259) | (5.263) | (5.510) |
| Exposure-robust (AKM), $t_{g_{\text{eff}}-1}$ | 12.734** | 8.967* | 13.953* | -15.949** | -18.510* | -16.361 |
| | (3.375) | (3.797) | (4.966) | (2.970) | (5.783) | (5.849) |

*Notes:* Daily minutes of market work; samples, controls, fixed effects, weights and (unless stated) commuting-zone–year clustering as in Table 4. Weekday diaries drop Saturdays and Sundays; the trend-augmented row adds six predetermined 1990 commuting-zone characteristics interacted with year effects; the lagged row instruments with the EU5 shifter measured three years before the diary; the IPW row reweights by the trimmed inverse probability of couple-with-young-children status (probit on demographics and the instrument). The inference rows keep the baseline point estimates and change only the standard errors: clustering by commuting zone, and the exposure-robust (AKM) procedure with $t$ critical values at $g_{\text{eff}} - 1$ degrees of freedom. Full outcome sets: Online Appendix Tables A18, A13, A5, A10, A29, A28. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

**Usual weekly hours.** Diaries record one day, so Table 6 re-estimates the baseline on respondents' usual weekly hours at all jobs, for most respondents the figure carried forward from the CPS interview months earlier. The female withdrawal appears here too, at 1.0 to 1.2 hours per week, consistent with the daily estimates; the male increase does not, as Section 2 predicts if men adjust on same-day realizations, extra shifts and overtime, that a usual-hours question smooths over. The day-level decomposition of Section 3.3 and the spouse-reported hours of Online Appendix Table A23 bracket the same asymmetry.

**Trade.** Import competition from China (Autor et al., 2013) hit the same regions over the same years. Conditioning on the commuting zone's China shock, instrumented with Chinese exports to other high-income economies, leaves the robot coefficients on market work and

Table 6: Effect of robot shock on hours worked per week (2SLS).

| | Males | | | Females | | |
|---|---|---|---|---|---|---|
| **Weekly hours worked** | (1) | (2) | (3) | (4) | (5) | (6) |
| Robots | 0.061 | 0.059 | -0.211 | -1.084** | -1.239*** | -1.023** |
| | (0.325) | (0.203) | (0.237) | (0.441) | (0.385) | (0.417) |
| Observations | 3,954 | 6,354 | 3,671 | 6,996 | 4,689 | 4,221 |
| R-squared | 0.219 | 0.146 | 0.199 | 0.237 | 0.189 | 0.206 |
| Spouse employed | ✓ | | ✓ | ✓ | | ✓ |
| Respondent employed | | ✓ | ✓ | | ✓ | ✓ |

*Notes:* Sample includes married or cohabiting couples with at least one spouse working and children aged 10 or less. Columns 1 and 4 capture the extensive margin of the respondent (spouse is employed). Columns 2 and 5 capture the extensive margin of the spouse (respondent is employed). Columns 3 and 6 restrict to dual-earner households, capturing only the intensive margin of labor supply. Control variables include commuting zone and state-by-year FE, household demographics, day/month/holiday. Standard errors clustered at the commuting-zone–year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

leisure unchanged on the 2003–2014 window where the trade measure is available, and the difference between those estimates and the headline ones comes from the window's sample, and hardly at all from the control (Online Appendix Tables A40 and A39). The trade shock itself produces a different household response, a sharp fall in men's market hours with a muted female response, consistent with an added-worker adjustment to an extensive-margin loss of manufacturing jobs, a different mechanism from the relative-wage channel of automation; Online Appendix Section F reports the comparison in full.

**Further checks.** Four further checks live in the Online Appendix. On the automotive concentration of the identifying variation, controlling for commuting-zone trends across quartiles of the 1990 automotive share leaves every estimate essentially unchanged (Table A9), while rebuilding exposure and instrument without the automotive component entirely leaves too weak an instrument for precise market-work estimates, though the childcare reallocation survives (Table A10), the shares-concentration disclosure of Section 3.1 seen from the other side. Beyond automotive: the male increase is 3–4 minutes larger among hourly-paid workers, as a flexible-hours adjustment implies (Table A20); the response is directionally concentrated where couples sit near earnings parity, significantly so on the leisure and paternal-childcare margins (Tables A21 and A22); and the norm-heterogeneity interactions of Table 7 keep their significance under commuting-zone clustering except on the broadest work margin (Table A30).

## 3.5 The mechanism: local gender norms

The causal estimates say what households do; this subsection asks why. The reallocation of Section 3.3 is at odds with what a frictionless household would do. Because the shock lowers both spouses' wages and the husband's by more (Section 3.2), such a household faces a fall in full income together with a rise in the wife's comparative advantage in the market. The income effect raises both spouses' market work, the added-worker response (Stephens Jr., 2002); the comparative-advantage effect reallocates market work toward the wife and childcare toward her husband. The level of the wife's hours is not pinned down, since her own wage falls and a bargaining weight that rises with her relative wage may raise her leisure (Chiappori, 1992; Blundell et al., 2005), but the direction of relative specialization is, and the evidence shows the opposite reallocation: market work moves toward the husband and childcare toward the wife. The argument proceeds by elimination, taking up in turn local demand, childcare prices, the direct displacement of women, and skill complementarity, before assembling the evidence for the remaining explanation, a within-couple friction activated by relative earnings. The bargaining alternative is estimated structurally in Section 4.4, where distribution-factor weights of increasing flexibility compete with the norm and fail to reproduce the reversal in fathers' childcare.

A generalized fall in local demand or a productivity spillover would depress market time for anyone in the exposed zones. Among singles, who face the same local environment without a within-couple margin, the effects on market work and leisure are instead null, and single mothers' childcare *falls* with exposure ($-8.3$ minutes), the opposite of partnered mothers and the sign flip a household mechanism predicts while a demand story does not (Online Appendix Table A24). If the withdrawal instead reflected childcare becoming unaffordable, prices would have to move enormously: measured childcare prices move by about 1% per unit of exposure, in offsetting directions across care types (Online Appendix Table A25), so that accounting for the female response would require a childcare-price elasticity near $-14$, two orders of magnitude beyond the consensus range.

Direct displacement of women is ruled out next, by four tests reported in full in Online Appendix Section E. Rebuilding exposure by sex from Census microdata shows that a commuting zone's exposure is, to a first approximation, the exposure of its men: robotized employment is male (automotive was 21 percent female in 1990), men's exposure alone reproduces every estimate, and the two sexes' exposures are too collinear to enter jointly (Online Appendix Tables A31–A32). Within women, the response sits where displacement cannot put it, among employed women *outside* manufacturing (94 percent of them), while women in robotized industries respond *less* (Online Appendix Table A33) and women whose own occupations carry a higher robot dose work *more* (Online Appendix Table A35). And when each spouse is assigned the predetermined exposure of their demographic cell, neither spouse's individual exposure moves women's market work, while men in exposed cells raise theirs (Online Appendix Table A34). The female withdrawal is therefore neither direct displacement nor a response to the husband's individual shock; it is a response to the *local* labor market moving against the male-breadwinner occupation mix, which is also what the couple-level first stage of Section 3.2 implies.

A skill-complementarity reading fails as well. If robots complemented high-skill labor, the female withdrawal would concentrate among lower-skill women, with no role for identity. Online Appendix Table A26 interacts the shock with the wife's education and rejects this reading: the market-work withdrawal is statistically indistinguishable between college- and non-college-educated wives, with the interaction small and insignificant in every column. The only margin that varies with education is childcare, where college wives raise their caregiving by somewhat less (an interaction of $-2.7$ minutes per robot in dual-earner households, Column 6), consistent with higher-income households substituting toward market-based childcare. The core of the reallocation holds across the skill distribution, as the norm channel predicts and a complementarity channel does not.

What remains must explain why a negative shock to the husband's wage lowers the wife's labor supply even as the household grows poorer, and the added-worker anomaly has exactly the structure a breadwinner norm implies. Automation lands hardest on male earnings in the lower-middle of the distribution, precisely where an adverse shock moves couples toward the threshold at which the wife earns as much as her husband (Bertrand et al., 2015), and the descriptive profiles of Section 2.1 kink at that threshold. If crossing it carries a social or psychological cost, an implicit tax on the wife's earnings once they threaten the male-breadwinner hierarchy, then the return to her work is taxed away faster than the income effect raises it, and her labor supply falls.

The first direct test uses spatial variation in the salience of the norm. I measure traditional gender culture with the Churches and Church Membership in the United States, 1990 survey of the Association of Statisticians of American Religious Bodies (ASARB); the 1990 baseline predates the robot wave, so a zone's religious composition cannot respond to the shocks under study. Following the literature that associates specific religious affiliations with traditional gender attitudes (Glass and Nath, 2006; Guiso et al., 2003), I compute the share of each county's population belonging to denominations that emphasize traditional family structures, among them the Church of Jesus Christ of Latter-day Saints (Mormons), the Southern Baptist Convention, the Assemblies of God, the Lutheran Church–Missouri Synod, and the Presbyterian Church in America, and I aggregate the shares to the commuting-zone level, normalizing by total population.

Traditional adherence concentrates in the Southeastern United States and the Intermountain West (Figure 3, panel (b)), a geography distinct from the automotive Midwest that carries the robot shock; the two variables correlate at $-0.18$ across zones. The weak correlation is central to the design: it means the interactions below can capture the mediating role of norms without merely repackaging the average effect of automation in high-manufacturing regions. Figure 4 separates the proxy from industrial structure. Residualizing both the 1990 gender employment gap and adherence on the local manufacturing share, zones with more traditional adherence show a significantly wider gender employment gap (unweighted across zones, as in Table 2, Panel A; robust $t = 2.8$). The association is descriptive, but it survives conditioning on the industrial base the shock loads on, precisely the property the design requires of the proxy.

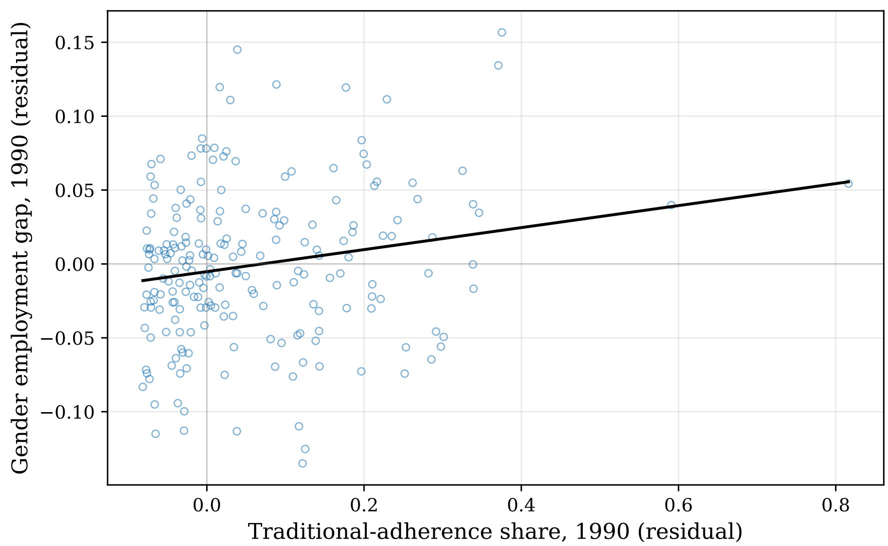


Figure 4: Added-variable plot of the 1990 gender employment gap (male minus female employment rate, ages 18–64) on the 1990 traditional-adherence share across commuting zones, both residualized on the 1990 manufacturing employment share. Each marker is one of the 209 commuting zones whose counties are identified in the 1990 five-percent Census sample; the line is the unweighted fit, the convention of Table 2, Panel A ($\hat{\beta} = 0.075$, robust $t = 2.8$).

I define a binary indicator $D_c^{High}$ equal to one if a commuting zone's share of traditional adherents is above the sample median, and estimate the augmented specification

$$Y_{ict} = \alpha + \beta_1 \text{Robots}_{ct} + \beta_2(\text{Robots}_{ct} \times D_c^{High}) + \eta' C_{ict} + \gamma_c + \delta_{st(c)\times t} + \varepsilon_{ict}, \tag{5}$$

where $\beta_1$ is the response in socially progressive regions and $\beta_2$ the additional response in conservative ones; both endogenous terms are instrumented with the corresponding interactions of the European instrument.

Table 7 shows a clear gender asymmetry in the interactions. For men (Columns 1–3), the market-work interaction is statistically insignificant: their response does not vary with local cultural priors. For women (Columns 4–6), the heterogeneity is substantial. In progressive commuting zones, an additional robot per 1,000 workers reduces women's daily market work by 17 minutes (Column 6); in conservative zones the reduction deepens by a further 34 minutes, so the total response in a high-exposure, conservative environment is nearly triple the progressive one. The withdrawn hours reappear at home: the childcare interaction is largest in the broad sample (+24 minutes, Column 4), and the leisure interactions are significant among employed and dual-earner wives (+15 and +21 minutes). The shock pushes households toward traditional specialization most strongly where the norm is culturally salient. The pattern does not depend on the binary split either: re-estimated with the *continuous*, z-scored adherence share, a one-standard-deviation increase in local traditionalism deepens women's market-work withdrawal by 30 to 37 additional minutes per robot (Columns 4–6, each significant at the ten-percent level; Online Appendix Table A27), so the heterogeneity scales smoothly with the intensity of local norms and does not simply switch on at an arbitrary cutoff. Both endogenous terms are strongly instrumented, with Sanderson–Windmeijer conditional first-stage $F$-statistics ranging from 324 to 476 for the level term and from 67 to 145 for the interaction across columns. As with the baseline estimates, the amplified coefficient is read within the interquartile range of the shock.

The asymmetry between the two spouses' responses is itself a prediction of the mechanism. The identity friction operates on the wife's side of the budget: as the husband's relative wage erodes and the couple moves toward parity, the wife's marginal return to market work falls while the husband's own return does not move. He adjusts only through the household's income, the added-worker motive that also operates on her, and through the reallocation of the hours she gives up, so his response is a second-order consequence of hers. The estimated model of Section 4 makes the point quantitatively: feeding the paper's calibrated erosion of the husband's wage through the benchmark, the wife's market-work share falls by 2.1 percent and the husband's rises by 0.8 percent, a male response about six-tenths of the female one in absolute value, with his childcare falling and hers rising, the same ordering and roughly the same ratio as the diary estimates. Without the wedge both signs flip: the husband substitutes away from market work as his own wage falls while the wife expands. Two features of labor supply reinforce the asymmetry. Married women's market hours, and mothers' of young children above all, respond to wages and to household income far more than men's, whose intensive-margin elasticities are close to zero and whose working time is tied to institutional full-time schedules (Blundell and MaCurdy, 1999; Keane, 2011; Blundell et al., 2016). And the margins on which the diary records the female response, a day without market work and a shorter working day (Online Appendix Table A17), are margins that part-time and flexible arrangements open to mothers far more than to fathers. A female response accompanied by a weaker, less precisely estimated male mirror image is therefore the pattern the norm predicts, and the one the data show.

The evidence lines up on one reading. The response is absent for singles, cannot be priced as a childcare-cost effect, is not displacement of women at any level of disaggregation, industry, occupation, or demographic cell, and deepens exactly where traditional norms are salient, although the norm's geography is orthogonal to the shock's. What remains is a within-couple friction that binds as couples approach earnings parity. The next section builds that friction into an estimable model.

# 4 A structural household model of time use

Section 3 establishes what happens and through which channel; this section asks how large the friction must be. I develop and estimate a semi-cooperative household model in which a breadwinner norm enters as a wedge on the wife's effective earnings. The model has four jobs: to rationalize the kinked time-use patterns of Section 2.1; to test the norm against every bargaining alternative on the one margin that separates them; to quantify the distortion, by comparing the estimated household against the same household with the wedge switched off; and to price the policy instruments that operate on the friction.

## 4.1 Environment and equilibrium

A household consists of two members, $i \in \{f, m\}$, each endowed with one unit of time, earning a market wage rate $w_i$, and jointly caring for $n$ children. Spouses play a two-stage game: in the first stage they cooperatively choose market hours; in the second, each chooses the time devoted to childcare non-cooperatively, taking market hours as given. Appendix A discusses the reasons for each modeling choice; this subsection states the model.

**Time budget.** Each individual faces the time constraint

$$h_i + l_i + (t_i + \tilde{t}_i)\, n \leq 1, \quad h_i \geq 0,\ l_i > 0,\ t_i \geq 0,\ \tilde{t}_i > 0, \tag{6}$$

where $h_i$ denotes market work time, $l_i$ leisure time, $t_i$ parental quality time invested per child, $\tilde{t}_i$ the fixed (non-quality) childcare time cost per child, and $n$ the number of children.

Table 7: Heterogeneity by local gender norms (IV estimates)

| | Males | | | Females | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| **Daily minutes work** | | | | | | |
| Trad × Robots | -8.250 | 16.430 | 0.229 | -26.982* | -29.814** | -33.686** |
| | (18.146) | (13.399) | (17.140) | (14.004) | (13.024) | (16.024) |
| Robots | 11.764* | 11.360** | 13.980** | -15.881*** | -19.354*** | -17.256*** |
| | (6.325) | (4.494) | (6.488) | (4.625) | (5.299) | (5.737) |
| Mean dep. var. | 260.2 | 281.7 | 278.6 | 134.3 | 218.8 | 216.5 |
| Observations | 4,100 | 6,593 | 3,812 | 7,194 | 4,891 | 4,400 |
| R-squared | 0.425 | 0.443 | 0.476 | 0.285 | 0.422 | 0.432 |
| **Daily minutes childcare** | | | | | | |
| Trad × Robots | 1.123 | 0.998 | -4.302 | 23.577** | 9.266 | 7.705 |
| | (7.456) | (6.937) | (6.970) | (9.190) | (5.753) | (6.736) |
| Robots | 0.986 | 2.611 | 1.265 | 6.519** | 9.949*** | 10.297*** |
| | (2.487) | (2.466) | (2.394) | (2.579) | (2.733) | (2.924) |
| Mean dep. var. | 85.6 | 80.8 | 82.5 | 148.0 | 118.6 | 120.6 |
| Observations | 4,100 | 6,593 | 3,812 | 7,194 | 4,891 | 4,400 |
| R-squared | 0.228 | 0.180 | 0.253 | 0.237 | 0.231 | 0.245 |
| **Daily minutes leisure** | | | | | | |
| Trad × Robots | -1.875 | -27.504*** | -5.131 | 5.166 | 15.066* | 21.138** |
| | (12.302) | (9.197) | (10.950) | (7.605) | (8.073) | (9.584) |
| Robots | -5.714 | -7.777** | -9.539** | 6.192** | 11.422*** | 10.308*** |
| | (4.410) | (3.750) | (3.705) | (2.947) | (2.910) | (3.477) |
| Mean dep. var. | 260.7 | 251.8 | 252.6 | 221.3 | 204.5 | 204.7 |
| Observations | 4,100 | 6,593 | 3,812 | 7,194 | 4,891 | 4,400 |
| R-squared | 0.289 | 0.259 | 0.313 | 0.208 | 0.305 | 0.313 |
| Commuting zones | 136 | 139 | 135 | 142 | 134 | 133 |
| Spouse employed | ✓ | | ✓ | ✓ | | ✓ |
| Respondent employed | | ✓ | ✓ | | ✓ | ✓ |

*Notes:* 2SLS estimates of robot exposure interacted with local gender norms; Trad equals one for commuting zones with an above-median 1990 share of traditional religious adherents. Samples, controls, fixed effects, weights and clustering as in Table 4. Both endogenous terms are instrumented; Sanderson–Windmeijer conditional first-stage $F$-statistics range from 324 to 476 (level) and from 67 to 145 (interaction) across columns. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

**Child quality.** The quality of children is produced with the two parents' time,

$$q = (1 + t_f)^{\alpha} \, (1 + t_m)^{1-\alpha}, \tag{7}$$

with $\alpha \in [0, 1]$ governing the relative effectiveness of maternal and paternal time. Only time inputs enter, as in Del Boca et al. (2014).

**Preferences and the identity wedge.** Individual $i$'s baseline utility is

$$u_i \;=\; \delta_i \, \log(c) \;+\; \mu_i \, \log(l_i) \;+\; \gamma_i \, \log(q \, n), \tag{8}$$

where $\delta_i$ is the weight on consumption; under logarithmic utility only the ratios of the preference weights are identified, so I normalize $\delta_f = \delta_m = 1$ throughout (the same $\delta_i$ appears in the closed-form solutions of Online Appendix J). Household consumable income is adjusted by the identity friction at the center of the paper, a wedge $\phi(s)$ on the wife's effective earnings, indexed to her share of *potential* wages:

$$c \;=\; \phi(s) \, w_f h_f \;+\; w_m h_m, \qquad s \;=\; \frac{w_f}{w_f + w_m}. \tag{9}$$

The wedge is a smooth, strictly decreasing and log-convex function of $s$,

$$\phi(s) \;=\; \exp\bigl[-\chi \, s^{\kappa}\bigr], \quad \chi > 0, \; \kappa > 1,$$

whose two parameters govern where along the wage-share distribution the friction sets in ($\kappa$) and how much consumable income it absorbs ($\chi$); both are estimated. Because $s$ is defined on potential wages, $\phi(s)$ is a constant in the household's time-allocation problem, and consumable income is linear in $(h_f, h_m)$.

**Stages.** The game is solved by backward induction.

*Stage 2 (Childcare):* Taking market hours $h_i$ as given from the first stage, each individual $i$ maximizes their private utility with respect to $t_i$:

$$v_i(t_i, t_{-i} \mid h) = \delta_i \log(c) \;+\; \mu_i \log\bigl(1 - h_i - (t_i + \tilde{t}_i) \, n\bigr) \;+\; \gamma_i \log\bigl((1 + t_f)^{\alpha}(1 + t_m)^{1-\alpha} n\bigr). \tag{10}$$

Because the child quality production function is Cobb-Douglas and utility is logarithmic, the objective function is additively separable in $t_f$ and $t_m$.[5] This stage results in four possible

[5]Specifically, the cross-partial derivative of spousal childcare time is zero ($\frac{\partial^2 v_i}{\partial t_i \partial t_{-i}} = 0$), implying orthogonal reaction functions. While this eliminates strategic substitutability, a structural "lack of commitment" friction

regimes: both, one, or neither spouse provides quality childcare time.

*Stage 1 (Market Work):* In the first stage, agents cooperatively choose market hours to maximize the household utility function, anticipating the Nash equilibrium outcomes of the second stage. The household objective is the weighted sum of individual utilities,

$$u_h = \theta\, u_f + (1-\theta)\, u_m,$$

where $\theta \in (0,1)$ denotes the bargaining power of the wife.

**Assumptions and equilibrium.** Under standard regularity conditions on the primitives, stated formally as Assumptions A1–A3 in Online Appendix Section I.1 together with Proposition A1, the two-stage game has a unique subgame-perfect equilibrium: the second-stage childcare game is diagonally strictly concave, so Rosen's theorem (Rosen, 1965) delivers a unique continuous Nash equilibrium for any market hours, and the induced first-stage household objective is strictly concave, so market hours are unique as well. Because the wedge is indexed to *potential* wages, $\phi(s)$ is a constant in the household's optimization and consumable income stays linear in hours, and the concavity argument rests on that linearity. Depending on the second-stage outcome, the solution falls into one of the twelve closed-form profiles reported in Online Appendix Section J.

Each ingredient of this environment serves one of the two regularities the model must reproduce, the corner solutions in daily childcare and the reversal of the profiles at parity. Non-cooperative childcare generates the corners: log utility and Cobb–Douglas quality make the spouses' care choices separable, so a parent stops providing quality time when market hours raise the shadow price of time above its private benefit. An additive utility cost would leave the consumption–leisure margin untouched; the wedge on the wife's *effective* wage bends her hours back at parity because it moves the marginal return to her work. Indexing it to *potential* wages keeps the wedge outside her control, the demanding "gender display" reading. And holding the bargaining weight fixed isolates the norm, with every wage-responsive weight estimated as a competitor in Section 4.4. Appendix A discusses each choice, and the mapping from assumptions to reproduced patterns, in full.

remains (Gobbi, 2018). In the semi-cooperative case, the FOC for spouse $f$ is $\frac{\alpha\gamma_f}{1+t_f} = \frac{\mu_f n}{l_f}$. In a fully cooperative (Pareto-efficient) benchmark maximizing $\theta v_f + (1-\theta)v_m$, the FOC becomes $\frac{\alpha\gamma_f}{1+t_f} + \frac{(1-\theta)}{\theta}\frac{\alpha\gamma_m}{1+t_f} = \frac{\mu_f n}{l_f}$. Because the social marginal benefit in the cooperative case includes the internalized positive externality to the partner ($\gamma_m$), the semi-cooperative Nash equilibrium leads to a structural under-provision of childcare ($t_i^{semi} < t_i^{coop}$). I estimate this fully-cooperative alternative as a competitor in the model comparison of Section 4.4: even with $\theta$ acting directly on care and free to kink and jump at parity, it cannot reproduce the paternal-care inverse-V (Table 10, Proposition 1).

## 4.2 Estimation

The model is estimated to reproduce the four time-use profiles of Section 2.1, market work and total childcare for each spouse, in both their levels and their kinked slopes across the wife's wage share.

**Calibrated parameters.** A first set of parameters is fixed a priori. The consumption weights are held at the normalization $\delta_f = \delta_m = 1$ of Section 4.1, which anchors the utility scale given that consumption quantities are unobserved. The child-quality productivity parameter $\alpha$ is fixed at 0.5, following Del Boca et al. (2014) and Gobbi (2018). Wages and the number of children are drawn from log-normal distributions parameterized with summary statistics from the ATUS sample (Online Appendix Table A41), with hourly wages deflated to constant-2000 dollars before the moments are computed.[6] The distributions for female wages, male wages, and the number of children are drawn independently, and the orthogonality serves an identification purpose: it ensures that the non-monotonic profiles generated by the simulation are driven strictly by the relative wage share activating the identity friction, with no help from mechanical demographic correlations, such as high-wage women supplying less childcare because they have fewer children. Accounting properly for the joint determination of fertility and human capital would require a dynamic life-cycle framework, which falls outside the scope of this static time-allocation model. The bargaining weight follows a logistic-normal distribution,

$$\theta_i \;=\; \frac{1}{1+\exp(-z_i)}, \qquad z_i \;\sim\; \mathcal{N}\big(\bar{\theta},\; \sigma^2\big),$$

which ensures $\theta_i \in (0,1)$. The remaining parameters,

$$\Theta = (\gamma_f,\, \gamma_m,\, \mu_f,\, \mu_m,\, \chi,\, \kappa,\, \tilde{t}_f,\, \tilde{t}_m,\, \bar{\theta},\, \sigma),$$

are estimated by the Method of Simulated Moments (MSM).

**Targeted moments.** The estimation targets the *shape* of the time-use profiles as well as their levels. For each of the four outcome blocks $b \in \{h_f,\, h_m,\, (t_f+\tilde{t}_f)n,\, (t_m+\tilde{t}_m)n\}$, I fit, identically on the data and on the simulation, a continuous two-segment line kinked at

[6] Deflation matters for the *dispersion* of the simulated wage draws: pooling nominal 2003–2023 wages injects cross-year inflation variance that spuriously widens the distribution of the relative wage share. Because a common annual deflator cancels in the within-household ratio $s = w_f/(w_f + w_m)$, deflation leaves the empirical (data) moments unchanged and only disciplines the simulation inputs. For the minority of non-working spouses, a potential hourly wage is imputed as the mean wage of employed individuals in the same education × year × sex × five-year-age × state cell, with a broader-cell fallback for thin cells.

earnings parity:

$$y_i \;=\; \ell_b \;+\; \beta_b^L\,(s_i - \tfrac{1}{2})\,\mathbf{1}\{s_i < \tfrac{1}{2}\} \;+\; \beta_b^R\,(s_i - \tfrac{1}{2})\,\mathbf{1}\{s_i \geq \tfrac{1}{2}\} \;+\; u_i, \tag{11}$$

estimated over the interior support $s_i = w_f/(w_f + w_m) \in (0.25, 0.75)$, where $\ell_b$ is the conditional level at parity and $(\beta_b^L, \beta_b^R)$ are the left- and right-of-parity slopes. The pair of slope signs encodes the time-use patterns directly: an inverse-V is $(\beta^L > 0, \beta^R < 0)$, a V is $(\beta^L < 0, \beta^R > 0)$. Stacking the level and the two slopes across the four blocks yields the moment vector

$$m \;=\; \big(\ell_b,\; \beta_b^L,\; \beta_b^R\big)_{b=1}^{4}, \qquad \dim(m) = 4 \times 3 = 12. \tag{12}$$

Leisure is excluded: because the time budget binds $(l_i = 1 - h_i - (t_i + \tilde{t}_i)n)$, it is the residual category and would be collinear with the targeted non-leisure uses. Fixing the knot of (11) at parity is a restriction the data can be asked about: re-estimating the two-segment model with a free knot on the same micro sample locates the reversal where the norm predicts, and imposing $k = 0.5$ is nearly costless in fit (Online Appendix Section G, Figure A6).

**The estimator.** Let $\hat{m}(\Theta)$ denote the simulated counterpart of $m$, computed by applying the same kinked regression (11) to the simulated households. The MSM estimator minimizes a weighted quadratic form with an additive bargaining restriction:

$$\hat{\Theta} \;=\; \arg\min_{\Theta}\; \big[m - \hat{m}(\Theta)\big]'\, W \,\big[m - \hat{m}(\Theta)\big] \;+\; \mathcal{P}(\Theta). \tag{13}$$

$W$ is diagonal. A level moment $\ell_b$ receives weight $w_\ell/\ell_b^2$ (a percentage-deviation loss), and each slope moment receives weight $\omega_b\, w_\beta / a_b^2$, where $a_b = |\beta_b^L| + |\beta_b^R| + c$ is the block's own slope amplitude (with a small floor $c$) so that a gently sloped profile is not numerically dominated by a steep one, and $\omega_b$ places extra weight on the male-childcare slopes. Following Duffie and Singleton (1993), this diagonal, self-normalizing matrix is used in place of the optimal inverse-variance weighting: the ATUS diary measures are zero-inflated and sparse in the tail bins, so the empirical moment covariance is noisy and ill-conditioned, and optimal weighting would mechanically down-weight the tail of the wage-share distribution that carries the identifying curvature. The chosen weights are $(w_\ell, w_\beta, \omega_{\text{M-care}}) = (6, 10, 2)$, selected by a grid search over the level–shape trade-off. Because this weighting is non-optimal by design, the Hansen over-identification test is a secondary check here; Online Appendix Section G reports the efficient-weight re-estimation and the test (the female-work reversal survives efficient weighting, the childcare curvature shrinks toward zero, and the test rejects, as it does for virtually any parsimonious static model at this sample size). The model is over-identified, with twelve moments for ten parameters, one of them disciplined externally.

Targeting the kinked slopes makes the curvature itself an explicit estimation target, and this choice sustains interior male quality care. Male childcare is $(t_m + \tilde{t}_m)\, n$, and its empirical inverse-V has an amplitude of only $\sim 0.01$; under an objective built on bin *levels* it would contribute almost nothing to the loss, and the cheapest way to match the flat $\approx 0.14$ level of male care is to drive men to the quality-care corner $t_m = 0$, where male care collapses to the constant $\tilde{t}_m\, n$, flat in $s$. With the slopes targeted, the wedge is identified from the reversal at parity instead of being optimized away. Under the self-normalizing weight the model reproduces simultaneously the three curvature reversals (the female-work inverse-V, the maternal-care V, and the paternal-care inverse-V), the rise of male market work above parity, and the four levels, with the male levels within 3% and the female levels within 15% (Table 9, Figure 6). The one feature the benchmark does not match in sign is the flat left-of-parity slope of *male* market work, which the model renders as a mild decline, the comparative-advantage response to a rising wife's wage; it is not one of the norm's identifying margins. The headline shapes, including the male-childcare inverse-V, also survive re-estimation under equal weights $(1, 1, 1)$, so they do not depend on the chosen weighting (Online Appendix Table A43).

**Identification.** Identification follows the standard mapping of moments to parameters. The *levels* $\ell_b$ identify the preference weights $(\gamma_i, \mu_i)$ and the fixed care costs $\tilde{t}_i$ that set how much time each spouse devotes to each activity; the *curvature* of the profiles, the reversal of the slopes at parity, identifies the wedge parameters $(\chi, \kappa)$, since only the norm bends a profile back at $s = 0.5$ (Section 4.4). A profiling exercise makes the identification of $(\chi, \kappa)$ explicit: fixing each wedge parameter on a grid and globally re-optimizing the other nine, the objective rises by roughly 18 SSR units as $\chi \to 0$ (from $\approx 5.7$ near the optimum to 23.8 at $\chi = 0$) and by about 12 as $\kappa \to 1$ (to 18.1), and the paternal-care inverse-V disappears everywhere at or below $\chi = 4$. The data therefore sharply reject the absence of a convex wedge, and any wedge much smaller than the estimated scale. Above the estimate the profile is shallower: fixing $\chi$ at 6 or 8 costs 1.2–1.3 SSR units and keeps all four shapes. The presence and the lower edge of the wedge's scale are thus pinned tightly, while its upper edge is identified only up to a range. The profiled near-optimal set (re-optimized SSR within $\approx 0.3$ of the profile floor) spans $\chi \in [4.6, 5.7]$ and $\kappa \in [2.0, 2.5]$, an implied parity-level tax of 62–68%, consistent with the bootstrap standard error on $\chi$ of 0.34. The sensitivity matrix of Andrews et al. (2017) confirms this reading: every free parameter draws at least 97% of its absolute sensitivity from the slope moments, and both wedge parameters load most heavily on the arms of male market work, the profile the wedge bends through the wife's withdrawal (Online Appendix Section G, Table A44). The female fixed care cost $\tilde{t}_f$ is estimated at 0.132, roughly 1.9 hours per child per day of non-quality care; its level is loosely pinned, and the shapes require only that it not be small (the profiling detail is in Online Appendix Section G).

One further parameter is weakly identified by the time-use shapes, and the same profiling makes this explicit. The mean female Pareto weight is shape-invariant: all four V and inverse-V signs are reproduced at every profiled point with implied $\mathbb{E}[\theta]$ between 0.44 and 0.55, and the objective is shallow across that range (SSR between 5.7 and 6.4 as the other parameters re-optimize), with no sharp optimum. Because the collective-household literature consistently places women's average bargaining weight below parity (Knowles, 2013; Friedberg and Webb, 2006), I impose this as an external-validity restriction instead of estimating the weight freely. The penalty $\mathcal{P}(\Theta) = w_\theta \max\{0,\ \mathbb{E}[\theta(\Theta)] - \bar{c}\}^2$ in (13) is one-sided: it is zero whenever the implied mean weight is already below the cap $\bar{c} = 0.47$, leaves the estimator free elsewhere, and targets no specific value. The time-use shapes do not depend on it.

**Estimates.** Table 8 reports the estimates with bootstrap standard errors. The identity-wedge scale is precisely estimated ($\chi = 4.67$, s.e. 0.34) and the implied mean female Pareto weight is significantly below parity ($\mathbb{E}[\theta] = 0.43$, s.e. 0.01), consistent with the collective-household literature. Figure 5 plots the estimated wedge itself, together with the envelope of the near-optimal profiled $(\chi, \kappa)$ pairs from the identification analysis above, and two features stand out when the wedge is expressed in economic units. First, the friction the reversal requires is large. At the benchmark, the household's consumable value of a dollar of the wife's earnings is 79 cents at $s = 0.25$, 53 cents at $s = 0.40$, and 35 cents at parity, and the implied tax averages 57% over the estimation sample's wage-share distribution. For a wife at parity earning the sample mean of \$758 per week (Table 1), the wedge is the resource equivalent of redirecting roughly \$489 per week away from consumable household income. These figures measure the as-if cost of the compensatory behaviors the wedge stands for, compensatory consumption and domestic friction, the "gender display" of Brines (1994); no cash outlay is involved. Overturning the added-worker response to an 18% wage shock requires a friction of this scale. Second, the figure shows what the profiles pin about the wedge: every near-optimal curve implies a mild wedge in the left tail (a tax of 16–25% at $s = 0.25$), a severe one in the right tail (92–94% at $s = 0.75$), and a parity-level tax between 62% and 68%, the identified set documented above. The final column of Table 8 re-estimates the model with the wedge removed ($\phi = 1$). Its fit is roughly four times worse (objective 23.7 versus 5.6), and the allocation it settles on is flat: with no wedge to bend the profiles, the estimator pushes both leisure weights and the bargaining dispersion to their bounds and still delivers monotone female work and childcare profiles with no reversal anywhere. This is the original intuition for the norm (Bertrand et al., 2015) in a different guise: absent an identity friction, no admissible configuration of preferences and bargaining reproduces what the profiles do at parity. The data favor the wedge over bargaining power, and Section 4.4 demonstrates this directly through the shapes.

Table 8: Structural parameter estimates

| Parameter | Description | Benchmark ($\phi \neq 1$) | No norm ($\phi = 1$) |
|---|---|---|---|
| $\gamma_f$ | Child-quality preference (wife) | 11.73 (0.12) | 10.62 (1.22) |
| $\gamma_m$ | Child-quality preference (husband) | 8.65 (1.08) | 11.86 (0.28) |
| $\mu_f$ | Leisure preference (wife) | 2.57 (0.27) | $5.00^{\ddagger}$ |
| $\mu_m$ | Leisure preference (husband) | 2.61 (0.56) | $5.00^{\ddagger}$ |
| $\chi$ | Identity-wedge scale | 4.67 (0.34) | $0^{\S}$ |
| $\kappa$ | Identity-wedge curvature | 2.17 (0.12) | $—^{\S}$ |
| $\tilde{t}_f$ | Fixed childcare cost (wife) | 0.132 (0.005) | 0.172 (0.008) |
| $\tilde{t}_m$ | Fixed childcare cost (husband) | 0.060 (0.003) | 0.074 (0.002) |
| $\bar{\theta}$ | Bargaining-weight mean | −0.91 (0.19) | −0.11 (0.07) |
| $\sigma$ | Bargaining-weight s.d. | 4.87 (0.92) | $8.00^{\ddagger}$ |
| $\mathbb{E}[\theta]$ | Implied mean female Pareto weight | 0.43 (0.01) | 0.50 (0.00) |
| Objective | Level+shape SSR (full $N$) | 5.6 | 23.7 |

*Notes:* Estimated by MSM on the 12 level-and-slope moments of (12). Standard errors (in parentheses) for *both* columns are from a household bootstrap ($B = 60$) that resamples the ATUS sample, recomputes the moments, and re-estimates. $\alpha = 0.5$ and the consumption weights ($\delta_f = \delta_m = 1$) are fixed throughout. Every benchmark parameter is interior to the estimation box (the box itself was widened once, on the dispersion and fixed-cost dimensions, when a preliminary estimate reached its edge; the reported estimate does not). The no-norm column re-estimates the model with the wedge switched off ($\chi = 0$), disciplined identically. §Both wedge parameters are excluded from the no-norm fit: $\chi$ is set to zero and $\kappa$ is then *unidentified* (under $\chi = 0$, $\phi(s) = \exp[-\chi s^{\kappa}] \equiv 1$ for any $\kappa$), so it is held fixed and not reported as an estimate. ‡Parameter at its bound in the no-norm fit; the wedge-free model pushes both leisure weights and the bargaining dispersion to their bounds and still erases every reversal, which is why its fit is roughly four times worse.

## 4.3 Model fit and validation

Table 9 contrasts the empirical averages for market work and childcare with the simulated moments generated under the benchmark parameterization. The simulation closely tracks the observed division of labor. Men devote 39% of their time budget to market work and 14% to childcare (against 38% and 14% in the data), while women allocate 24% to market work and 34% to childcare (against 22% and 30%): the two male levels are matched within three percent and the two female levels within fifteen percent. As Figure 6 shows, the model also reproduces the three V and inverse-V reversals, in female market work, maternal childcare, and paternal childcare, together with the rise of male market work once the wife out-earns her husband. The one slope it misses is the left arm of male market work: flat in the data,

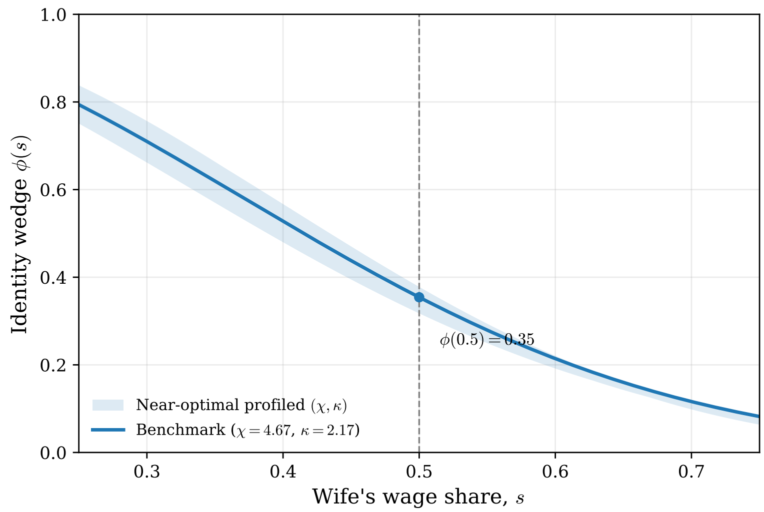


Figure 5: The estimated identity wedge $\phi(s) = \exp[-\chi s^{\kappa}]$ over the estimation support. The solid line is the benchmark estimate ($\chi = 4.67$, $\kappa = 2.17$); the shaded band is the envelope of the near-optimal profiled $(\chi, \kappa)$ pairs (re-optimized SSR within $\approx 0.3$ of the profile floor) and is an identified-set display; it should not be read as a sampling interval. Every near-optimal curve implies a mild wedge in the left tail and a severe one in the right tail; the level at parity (benchmark $\phi(0.5) = 0.35$) is pinned to $[0.32, 0.38]$, and the profiles reject any wedge with $\chi \leq 4$, so they identify the wedge's presence and minimum scale, beyond its convexity alone.

mildly declining in the model, where comparative advantage alone would move market work toward the wife.

Figure 6 overlays the simulated and empirical profiles for market work and childcare across the wife's wage share. For each block both data and model are summarized by the same object, a two-segment line joined at parity (the data fit in black, the model fit in blue), so that the model curve is read directly against the black empirical curve and no single noisy bin distorts the comparison. Female market work traces its inverse-V, expanding toward parity and contracting once the wife crosses the threshold; maternal childcare traces the mirror-image V; and male market work rises above parity (the model adds a mild decline below parity that the data, flat there, do not show). The gentle inverse-V in paternal childcare comes through as well, with left and right slopes $(+0.022, -0.032)$ that track the data point estimates $(+0.013, -0.035)$. The empirical paternal-care slopes are individually imprecise (left $t \approx 0.5$, right $t \approx -1.0$), so the data establish the *shape* of this profile far less sharply than the female-work and maternal-care reversals, whose arms are all significant; the profile's evidential weight comes from the model comparison of Section 4.4, where it is the one shape that only the wedge reproduces while every bargaining alternative gets its sign wrong. A household block bootstrap of the simulated profile places the inverse-V in 100% of resamples, so the model feature is not a simulation artifact. The estimated identity penalty is thus a sufficient friction to drive the kinked behavioral adjustments isolated in the reduced-form analysis.

**External validation.** Two objects the estimation never targets check the model out of sample. The first is the extensive margin. In the data, the probability that the wife does any market work on the diary day traces an inverse-V across $s$: the withdrawal of Section 2.1 operates on participation as well as hours. The benchmark reproduces the shape and, untargeted, essentially the level (a female participation rate of 0.40 against the diary-day rate of 0.38). Both no-norm alternatives are instead monotone: the wage-rising bargaining weight has participation rising in $s$ throughout, and the wedge-off re-estimate has the wife's participation rising and the husband's falling, with no reversal at parity (Online Appendix Figure A7). The model's single participation margin cannot separate labor-force attachment (0.68 for women) from the diary-day rate, so the informative comparison is the shape of the profile. The second is the one *within-couple* profile the survey carries: the spouse's usual weekly hours, reported at the ATUS interview, in the same household whose $s$ the estimation bins. Mean-normalized and fit with the same two-segment line, the wife's hours as reported by her husband trace a sharply significant inverse-V ($+0.86$ and $-0.53$ per unit of $s$, standard errors 0.12 and 0.15), and the husband's hours as reported by his wife rise significantly above parity ($+0.21$). The benchmark reproduces both, untargeted ($+0.63$ and $-0.97$; $+0.32$), with its one recurring miss, a mild decline of male hours below parity where the data are flat, repeated here. Re-estimating with these four spouse-hours slopes *added* to the moment vector moves the wedge little ($\chi = 4.44$ against 4.67) and preserves every reproduced shape, so the benchmark is retained and the spouse profiles are reported as validation. The within-couple reallocation the model implies is checked against within-couple data; nothing about it is assumed.

## 4.4 Norm versus bargaining power

A natural concern is whether the identity wedge is truly necessary, or whether the same non-monotonic profiles could arise from a standard collective channel in which the wife's Pareto weight rises with her relative wage (Chiappori, 1992; Browning and Chiappori, 1998). I therefore re-estimate four no-norm alternatives under the identical objective (13), in increasing order of flexibility. The first is the $\phi = 1$ model of Table 8, in which bargaining is constant in $s$. The second is a **monotone** weight rising linearly in $s$. The third is a **limited-commitment** weight that may kink *and* jump at parity, the shape a binding participation constraint produces in limited-commitment collective models (Mazzocco, 2007; Voena, 2015; Foerster, 2025). The fourth is the demanding case: because the first three inherit the semi-cooperative timing, in which the weight never enters the childcare first-order conditions, I add a **fully-**

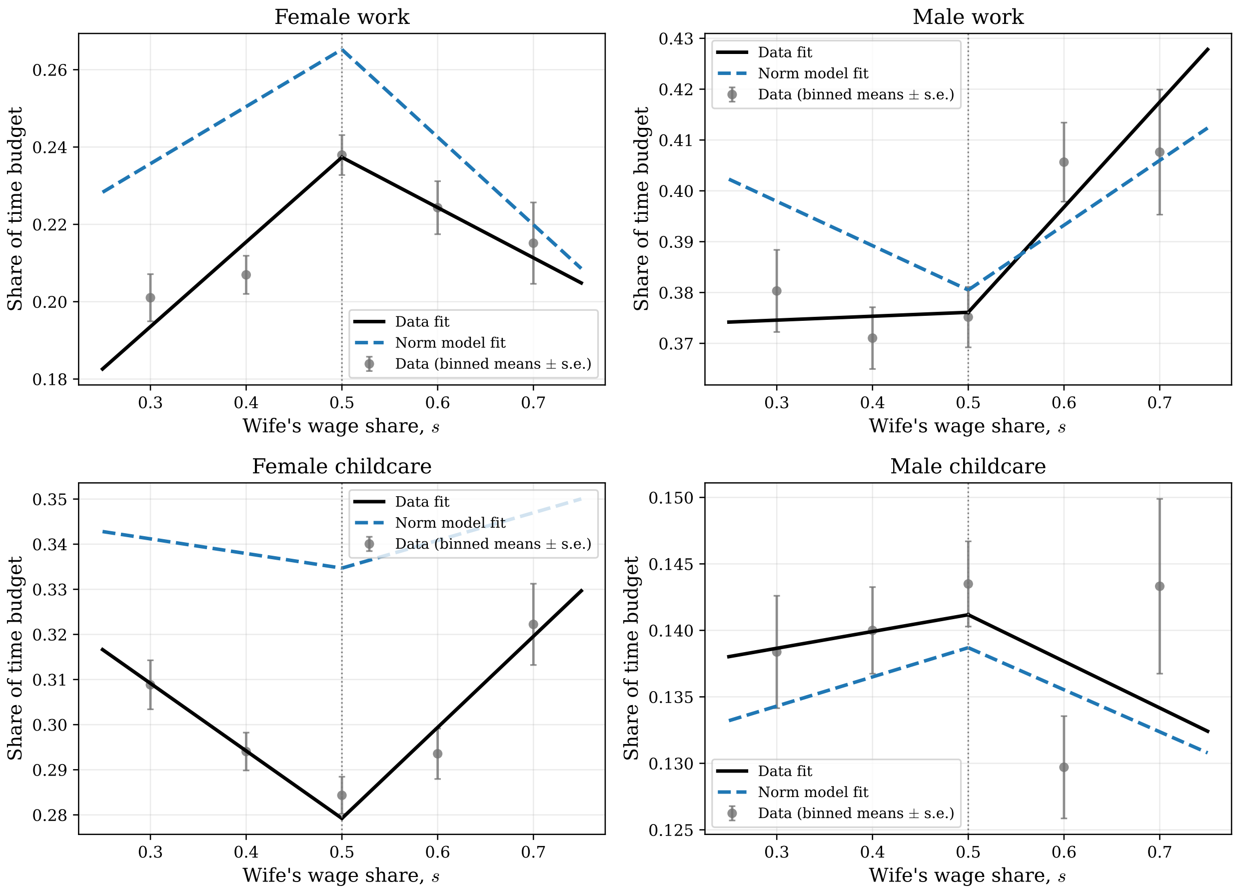


Figure 6: Fit of the benchmark norm model to the data, across the wife's wage share $s = w_f/(w_f + w_m)$. Grey markers are binned data means (± one standard error); the black line is the two-segment data fit and the blue dashed line is the benchmark norm model, both kinked at parity $s = 0.5$. The norm model is shown in blue throughout, matching the same model in Figure 7. This figure assesses the benchmark fit; the mechanism comparison against the bargaining alternatives, using the *same* benchmark norm curve, is in Figure 7.

Table 9: Comparison of model and data averages (benchmark model)

| | | | Data | |
|---|---|---|---|---|
| Sex | Moment | Model | Diary | Employed |
| **Internal moments (targeted levels)** | | | | |
| Male | Work | 0.393 | 0.383 | |
| | Childcare | 0.136 | 0.139 | |
| Female | Work | 0.244 | 0.218 | |
| | Childcare | 0.340 | 0.296 | |
| **External moments (untargeted participation)** | | | | |
| Male | Participation $P(h > 0)$ | 0.99 | 0.61 | 0.97 |
| Female | Participation $P(h > 0)$ | 0.40 | 0.38 | 0.68 |

*Notes:* Targeted levels are shares of the discretionary budget. Participation is untargeted; "Diary" is the rate of any market work on the diary day, "Employed" is the labor-force employment rate. The static model has one participation margin and sits near one data concept for each sex (see text); the informative comparison is the participation *profile* of Online Appendix Figure A7.

**cooperative** model in which the household chooses all four time uses jointly, so the same kinked-and-jumping weight acts on childcare *directly* (Online Appendix Section I.2).[7]

Table 10 reports the outcome, and Figure 7 draws it. Within the semi-cooperative timing the norm fits best (5.5 against 6.1, 8.1 and 23.7), and it is the only model that reproduces all four shapes. The constant weight leaves every profile flat or monotone. The monotone and limited-commitment weights bend female *market* work, since any $s$-varying weight can produce that inverse-V through the wife's first-stage hours, but they leave the childcare margins monotone or flat. The fully-cooperative weight is the demanding case: acting on childcare directly, it recovers the maternal-care V as well and edges the norm on the raw objective (5.2 against 5.5), so neither the maternal reversal nor overall fit alone identifies the norm. What no bargaining weight reproduces, including this one, is the *paternal*-care inverse-V. That shape is the signature, and the identification of the norm rests on the joint pattern of the reversals.

**Proposition 1** (Childcare-ratio invariance)**.** *In the interior regime of the fully-cooperative*

[7]The norm and the first three alternatives are nested in the same solved semi-cooperative model; the fully-cooperative competitor is a separate closed-form re-solution of the joint program, estimated under the identical objective. The limited-commitment and fully-cooperative weights are estimated without the bargaining cap, and the conclusion is robust to the functional form of the monotone weight (linear, log-wage-ratio and logit fit similarly, at 8.1–8.4, and none bends the childcare margins).

*program, the ratio of the two spouses' quality-care inputs is*

$$\frac{1+t_f}{1+t_m} = \frac{\alpha}{1-\alpha} \cdot \frac{w_m}{\phi(s)\, w_f},$$

*independent of the Pareto weight $\theta$. A change in $\theta$ scales both spouses' interior childcare by a common factor and cannot alter their relative allocation; only the identity wedge $\phi(s)$, which enters the wife's effective wage, moves the ratio.*

Proposition 1, proved in Online Appendix Section I.2, gives the structural reason that no amount of flexibility overcomes. Within any interior household a Pareto weight redistributes resources but leaves relative prices, and with them the care ratio, unchanged; only the wedge, by entering the wife's effective wage, changes the relative price of her time and bends the two care margins in *opposite* directions at parity. A wage-share–dependent $\theta(s)$ can still shift the population care *profiles* through the composition of households moving in and out of the work corner as $s$ rises, the channel through which the fully-cooperative competitor recovers the maternal-care V, but nothing in a Pareto weight bends fathers' care in the direction opposite to mothers'. This is also why disciplining the bargaining weight in the benchmark is innocuous: bargaining power, however flexible and however cooperative, is not the source of the childcare reversal.

Table 10: Which mechanism reproduces the V / inverse-V profiles?

| Model | SSR | Female work | Maternal care | Paternal care |
|---|---|---|---|---|
| Data | — | inverse-V | V | inverse-V |
| Identity norm $\phi(s)$ | 5.5 | inverse-V ✓ | V ✓ | inverse-V ✓ |
| No norm: constant $\theta$ | 23.7 | monotone | flat | flat |
| No norm: monotone $\theta(s)$ | 8.1 | inverse-V ✓ | monotone | flat |
| No norm: limited-commitment $\theta(s)$ | 6.1 | inverse-V ✓ | monotone | flat |
| No norm: fully-cooperative $\theta(s)$ | 5.2 | inverse-V ✓ | V ✓ | flat |

*Notes:* SSR is the full-$N$ level-plus-shape objective; ✓ marks a profile whose shape (the pair of slope signs of a line kinked at parity) matches the data. "Constant" fixes the bargaining weight $\theta$ in $s$; "monotone" lets $\theta$ rise linearly in $s$; "limited-commitment" lets $\theta$ kink *and* jump at parity; "fully-cooperative" additionally admits $\theta$ into the childcare first-order conditions by re-solving the joint program (Online Appendix Section I.2). The last two are estimated without the bargaining cap. Female market work does not discriminate, since every $s$-varying weight reproduces its inverse-V, and the fully-cooperative weight also recovers the maternal-care V and edges the norm on the raw objective; only the identity wedge reproduces the *paternal*-care inverse-V, which identifies the norm (Proposition 1).

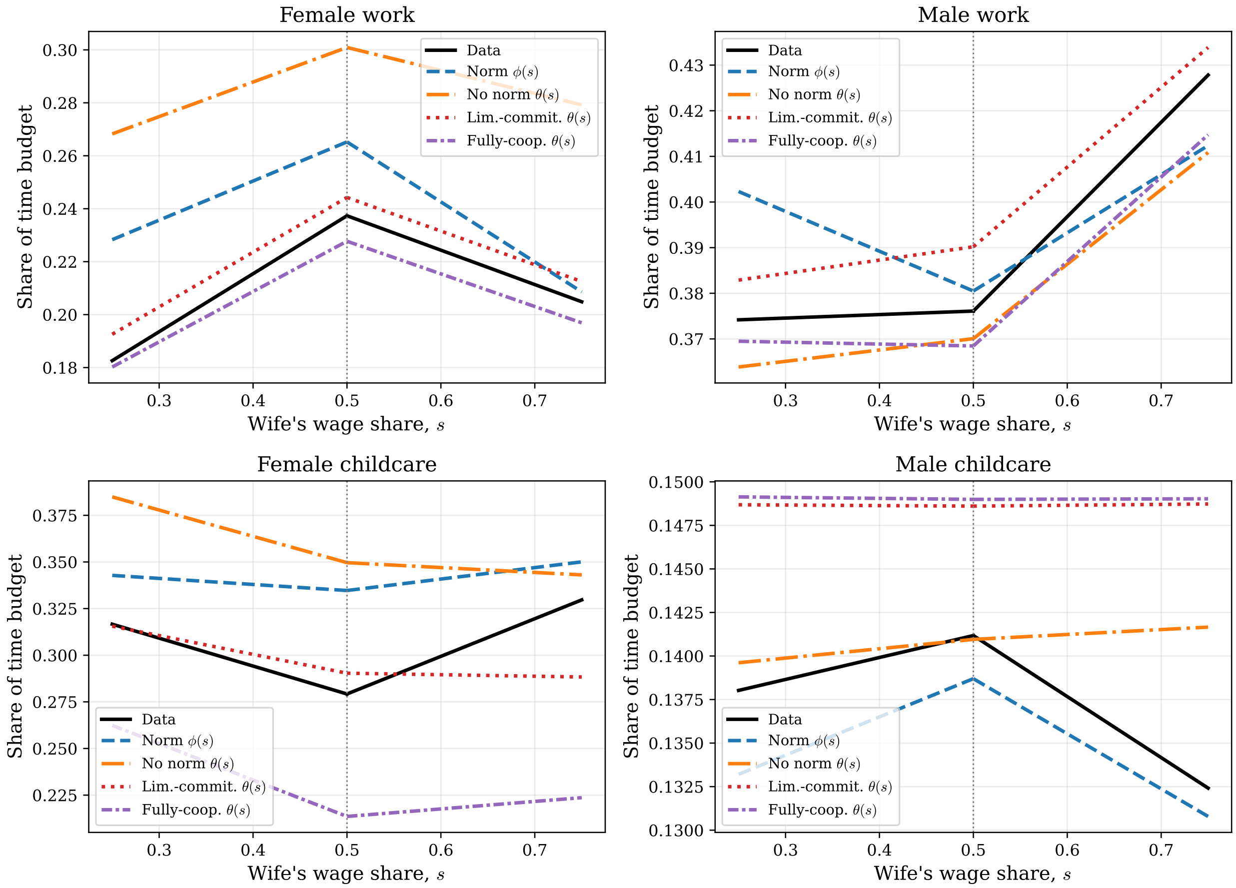


Figure 7: Two-segment fits kinked at parity: data (black), identity norm (blue, the *same* benchmark estimate as Figure 6), monotone bargaining (orange), limited-commitment bargaining with a kink and jump at parity (red), and the fully-cooperative weight that additionally bargains over childcare directly (purple); the bargaining models are re-estimated on the same moments. Every threshold weight reproduces the inverse-V in female *work* (top-left), and the fully-cooperative weight also recovers the maternal-childcare V (bottom-left, purple), which the semi-cooperative weights flatten. But only the norm reproduces the *paternal*-childcare inverse-V (bottom-right): every bargaining weight, including the fully-cooperative one, leaves fathers' care flat. The reversal in fathers' care separates the norm from its competitors; overall fit does not (Proposition 1).

## 4.5 Counterfactuals and policy experiments

The counterfactual feeds the paper's calibrated shock through the estimated model with the wedge switched on and off. Empirically, the robot wave compressed the gender wage gap mainly by eroding men's earnings in the lower-middle of the distribution (Acemoglu and Restrepo, 2020; Cortes et al., 2024); I therefore cut the potential wage of bottom-tercile husbands by 18.3%, three robots per thousand workers times the three-year wage estimate of Online Appendix Table A14, and follow the gaps among the shocked households across the range of multipliers in Figure 8. Any complementarity gains for high-skill women (Cortes et al., 2024) would raise $s$ further and amplify the contrast, so shocking only the husband is the conservative design. Because the same estimated parameters are used with the wedge on and off ($\chi \to 0$), the contrast isolates the wedge; no preference parameter is re-estimated along the way.

The two regimes answer the shock in opposite directions; the paths below run from no shock to a twenty-percent cut in bottom-tercile husbands' potential wages, marginally deeper than the calibrated 18.3%. Under the norm the wife *withdraws* as her relative wage rises ($h_f$: 0.234 $\to$ 0.229), so the work gap $h_m - h_f$ widens (0.164 $\to$ 0.173) and the care gap $t_f - t_m$ edges up (0.208 $\to$ 0.210), the reduced-form pattern of Section 3.3. With the wedge off, the standard logic operates: the wife raises her hours (0.356 $\to$ 0.366), out-works her lower-wage husband (work gap $-0.071 \to -0.096$), and the care gap, most of which the asymmetric fixed care costs generate in both regimes, narrows (0.161 $\to$ 0.157). The wedge therefore *reverses the sign* of the work gap, the data's positive gap of 0.17 against the efficient allocation's negative one, and widens the care gap beyond what the fixed costs imply; the work distortion between the two regimes is 0.24 of the time budget at baseline and 0.27 at the deepest exposure (Figure 8), the care distortion 0.05 to 0.06. Neither the disciplined Pareto weight nor sampling uncertainty drives the contrast. Sweeping $\mathbb{E}[\theta]$ across $[0.34, 0.64]$ leaves the norm's work gap positive throughout, the efficient gap negative everywhere below $\mathbb{E}[\theta] \approx 0.47$ (the benchmark is 0.43), and the distortion between 0.18 and 0.28. And across the $B = 60$ bootstrap re-estimates, the sign reversal and the added-worker contrast obtain in every draw, with a 90% interval of $[0.18, 0.26]$ for the baseline work distortion.

Two clarifications bound what the counterfactual claims. First, the large numbers are *level* contrasts between regimes. The model's own response to the shock is small: the 18.3% cut moves the average shocked couple's wage share by +0.04, pushes 8.7% across parity, and lowers the wife's market work by 2.1% (crossers $-3.4\%$, stayers $-2.0\%$, about one wife in a hundred leaving work), against the 37% contraction the reduced form implies for the same shock. No crossing rate could close that gap, so the within-couple relative-wage margin cannot carry the reduced form. The 2SLS estimate is instead the *total* local-labor-market treatment:

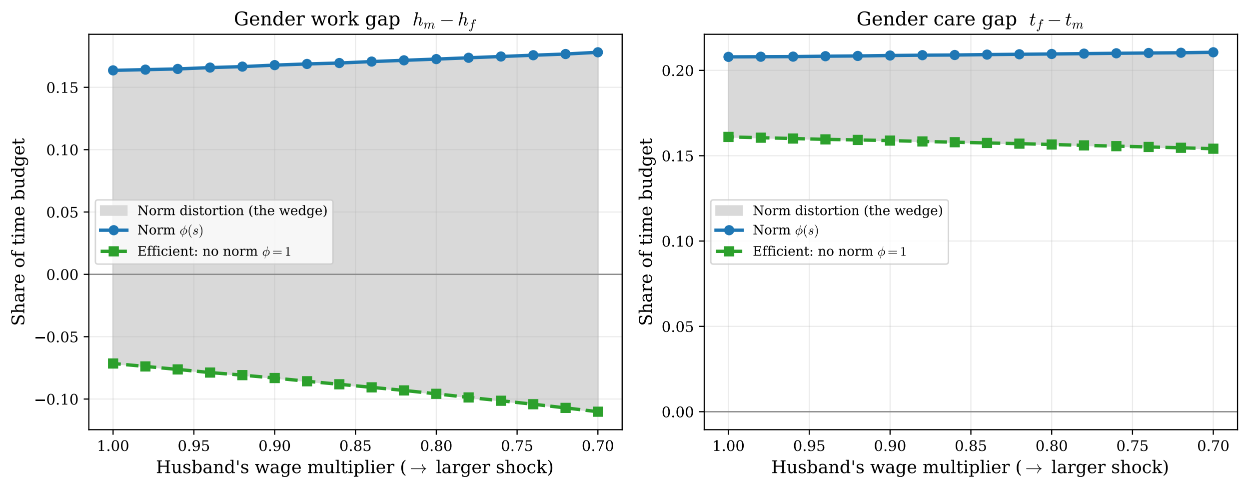


Figure 8: Simulated time-use gaps under negative shocks to bottom-tercile male wages: the benchmark (norm, blue) versus the identity-neutral allocation (green; the *same* estimated parameters with the wedge switched off, $\chi = 0$). The two panels are the gender work gap $h_m - h_f$ and the gender care gap $t_f - t_m$, plotted on a *common* vertical axis (share of the time budget). Moving right along the $x$-axis is a larger automation shock and a higher wife's wage share. For the work gap the norm stays *positive* while the efficient allocation is *negative*—the wife out-works the husband (the added-worker effect)—and grows more negative; the shaded band is the norm's distortion relative to the efficient allocation, which widens as exposure deepens.

direct labor demand, general-equilibrium adjustment, and the local activation of norms that the heterogeneity of Section 3.5 documents and the model, which holds $\chi$ fixed, does not capture. The structural contrast isolates something different, the norm's level distortion at any given wage distribution. Read as an elasticity, the reduced form is itself far outside what substitution can deliver (an implied cross-wage elasticity of female labor supply near +2, an order of magnitude above the own-wage range of 0.2–0.6 and opposite in sign to the insurance prediction, Chetty et al., 2011); a powerful friction is required, and the estimated wedge supplies one. Second, the model reproduces the empirical ordering of margins, with the work response exceeding the childcare response, as it must when market hours are the contractible first-stage margin and childcare adjusts only through the second-stage corner; the simulated and estimated responses agree in sign and ordering, and the quantitative gap comes from the shock channel and leaves the mechanism intact.

The estimated household can also price the policy instruments most often proposed for regions in automation's path, and three of them enter the model exactly. A credit at rate $\sigma$ on the wife's earnings raises her effective wage from $\phi(s)\, w_f$ to $(\phi(s) + \sigma)\, w_f$: a tax instrument moves take-home pay while both potential wages, and with them the reference share $s$, are unchanged.[8] The credit stands for the relief that individual taxation would grant

[8]The government's dollars are taken to reach the household intact, since the burn the wedge stands for is indexed to the wife's gross earning power; Online Appendix Section H reports the alternative convention, in which the display friction consumes the credit at the same rate as her market earnings, with the same

a secondary earner who now faces her household's marginal rate at her first dollar (Bick and Fuchs-Schündeln, 2018), and its mirror image, a tax at rate $\tau$ on the same earnings, is the phase-out that family-based transfers impose: the EITC withdraws 21 cents per dollar of a married secondary earner's income over its phase-out range (Eissa and Hoynes, 2004). Wage insurance at replacement rate $r$ pays the displaced husband $r$ of his hourly wage loss, the design proposed by Kletzer and Litan (2001) and implemented for older displaced workers, at $r = \frac{1}{2}$, as Reemployment Trade Adjustment Assistance; because the top-up is earned income it enters his potential wage, so this is the one instrument that acts on the share $s$ itself. Public childcare, finally, enters as provision in time, the only form a model with purely parental care can price: slots covering a share $\rho$ of the fixed custodial load lower the care costs to $(1-\rho)\,\tilde{t}_i$, and the experiment is accordingly silent on care prices, program take-up, and the quality of purchased care (Online Appendix Section H discusses the mapping). Each instrument is fed to the exposed couples in the post-shock world of Figure 8, with the wedge on and off and every parameter held at its estimate.

Each experiment is evaluated by the fiscal flow it generates, in constant-2000 dollars per exposed couple per week, against the female market hours it moves, so that a dollars-per-hour price ranks instruments and regimes alike; the outlay at which the exposed wives' hours return to their pre-shock level prices the repair of the shock's own damage. That damage is half an hour per week on the within-couple margin, so restoration is a weak hurdle, and the informative object is the price of an hour and how the wedge moves it. Table 11 collects both regimes at the headline doses; Figure 9 traces the two tax instruments across rates.

The credit is the cheap instrument, and the wedge is the reason. At $\sigma = 0.10$ the exposed wife adds 7.0 hours of weekly market work under the norm and 0.4 in the efficient household, at outlays of \$52 and \$65: an hour costs \$7 where the wedge binds and \$162 where it does not, and repairing the shock's damage requires a credit of half a cent on the dollar, \$2.40 per couple per week. Two features of the estimated household produce the asymmetry, and both work through $\phi$. First, the wedge suppresses the earnings on which the credit is paid, so the subsidy bill starts from a base her withdrawal has already shrunk. Second, over the exposed group's range of $s$ the effective wage is at most 35 cents on her dollar and falls toward zero in the right tail, so ten additional cents raise her price of time by 30 percent at parity and by multiples of it beyond, and the response runs through participation, which rises from 0.39 to 0.50. The ceiling of this logic is a credit indexed to the couple's own share, $\sigma(s) = 1 - \phi(s)$, which reproduces the efficient allocation outright at about \$560 per exposed couple per week, the as-if burn of Section 4.2 rendered as a fiscal bill and larger here than at parity because the exposed couples sit where the wedge is severe. The same arithmetic runs in reverse for

qualitative conclusions.

Table 11: The fiscal price of a female market hour under the policy instruments

| | $\Delta$ female hours/week | | Outlay | Dollars per hour | |
|---|---|---|---|---|---|
| Instrument (dose) | Norm | Efficient | (\$/week) | Norm | Efficient |
| Secondary-earner credit ($\sigma = 0.10$) | +7.0 | +0.4 | 52 | 7 | 162 |
| Wage insurance ($r = 1$) | +0.5 | −0.9 | 68 | 134 | 55 |
| Childcare provision ($\rho = \frac{1}{2}$) | +4.1 | +7.5 | 151 | 36 | 20 |
| EITC phase-out ($\tau = 0.21$) | −2.4 | −1.1 | −10 | 4 | 124 |

*Notes:* Exposed couples (bottom-tercile husbands) in the post-shock world, with the wedge on (norm) and off (efficient) at the same estimated parameters; instruments as defined in the text, credits and taxes under the gross-earnings incidence. Outlay is the fiscal flow per exposed couple per week in constant-2000 dollars under the norm (the efficient-regime flows are \$65, \$49, \$151, and \$131); a negative entry is revenue. Dollars per hour is the absolute flow per weekly hour of female market work the instrument moves; for insurance in the efficient regime the hours are removed, and for the phase-out the entry is revenue collected per hour destroyed. Childcare provision is priced at the childcare workers' median wage of \$7.20 per hour (CPS, 2000), one caregiver per child; group staffing divides the outlay by the staff ratio. Every sign pattern in the table obtains in all 60 bootstrap re-estimates; 90-percent bands are in Online Appendix Table A45.

the phase-out: stacking $\tau = 0.21$ on a parity wedge of 0.35 leaves the wife 14 cents on the dollar, and the tax destroys a weekly hour of her work for every \$4 of revenue under the norm against \$124 in the efficient household, an excess burden 30 to 78 times larger depending on the incidence convention (Online Appendix Section H).

Wage insurance operates on the trigger. Full replacement returns the exposed couples to their pre-shock wage distribution by construction, and with it the 8.7 percent of couples the shock pushed across parity (the wage path retraces Figure 8 in reverse; the policy experiment adds the ledger). Under the norm her hours rise with $r$, by half an hour per week at full replacement, because the penalty on her earnings recedes as $s$ falls back; in the efficient household the same policy lowers her hours by 0.9, the added-worker response dismantled as his wage is made whole. Insurance therefore protects female labor supply only where the norm binds, and it is the expensive protection: the program pays every displaced husband on his full hours, so a weekly female hour costs \$134 under the norm against the credit's \$7, on an outlay of \$68 per couple, 18 percent of the restored wage bill.

Childcare provision moves hours in both regimes, and the wedge decides where the freed time goes. Covering half the fixed custodial load frees 14 hours of parental time per week for the mean exposed couple, 10 of them the mother's. The efficient household sends 52 percent of the freed hours to market work; the norm household sends 29 percent and routes the balance to leisure and quality care, so the same slots yield 7.5 additional weekly hours of female work in one regime and 4.1 in the other, \$20 against \$36 per hour at a caregiver's wage. The care gap narrows under both regimes, from 0.21 to 0.16 of the time budget under the

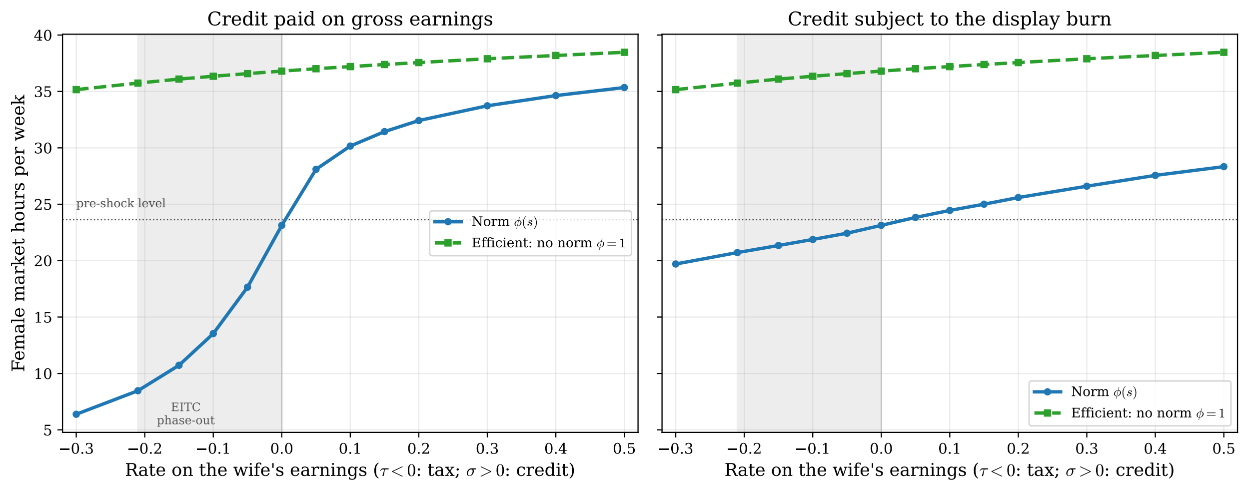


Figure 9: Female market hours of the exposed wives across the tax instruments: a rate on the wife's earnings, negative for the EITC-style phase-out $\tau$ and positive for the secondary-earner credit $\sigma$, under the benchmark norm (blue) and the identity-neutral allocation (green; the same estimated parameters with the wedge off). The left panel pays the credit on her gross earnings; the right panel subjects it to the display burn, the alternative incidence convention of Online Appendix Section H. The shaded band marks the EITC phase-out range and the dotted line the pre-shock level of female hours. Under the norm the same rate moves hours several times as far as in the efficient household on either side of zero: the wedge makes female hours cheap for the tax code to move in either direction.

norm at half coverage, and closes entirely under mother-only incidence at near-full coverage (Online Appendix Section H); provision is the one instrument among the three that moves the care margin whether or not the norm binds. The same margin bears on the cross-country childcare evidence: expansions moved maternal employment strongly in Quebec and Germany (Baker et al., 2008; Bauernschuster and Schlotter, 2015) and hardly at all in Norway (Havnes and Mogstad, 2011), and while the Norwegian result is attributed to crowd-out of informal care, a margin this model lacks, the counterfactual adds a second candidate: where the norm binds, freed time flows to leisure and quality care because the wedge has cut the market return to cents on the dollar.

The same parameters can be taken to the policy experiments in the norms literature. Lippmann et al. (2020) find the out-earning aversion absent where institutions eroded the norm, and Ichino et al. (2026) find childcare responding to the father's tax incentives only in couples from traditional-norm countries. Closest to the model's experiment, Grönqvist et al. (2026) find that a Swedish pay reform raising teachers' wages by about 15 percent narrowed the couple's childcare gap by 12 percent, with mothers doing less. Fed the same own-wage rise, the estimated model moves in the same direction but far less (female work +2 percent below parity; care gap −0.6 percent), and the wedge mutes it: with $\phi = 1$ the same raise moves work twice and care three times as much. An own-wage rise operates where the wedge is gentle, a husband's-wage fall where it steepens, which is why the first expands female work in both regimes while the second splits them; the U.S. model's muted response relative to Sweden is consistent with the weaker traditional norms of that setting and with Grönqvist et al.'s own finding that progressive couples adjust more.

# 5 Conclusion

This paper has examined how industrial automation reshapes the allocation of time within households. Linking American time diaries to commuting-zone robot exposure, with exposure instrumented by European adoption, I document a systematic within-couple reallocation: each additional robot per thousand workers leads women to withdraw close to two hours per week from market work, about 16 minutes per day, reallocated toward childcare and leisure, while men's diary minutes move in the opposite direction, more weakly.

A structural household model rationalizes these patterns through an identity wedge that penalizes the household as the wife's earning power approaches her husband's. Estimated by simulated method of moments on the kinked time-use profiles, the model reproduces the observed reversals, and the model comparison shows that only the wedge bends the two parents' childcare in opposite directions at parity: no bargaining weight does so, not even one that bargains cooperatively over childcare and is free to kink and jump at the threshold. In the counterfactual, the norm reverses the efficient response to the shock. An unconstrained household would answer the erosion of the husband's wage by having the higher-wage wife expand her market hours and out-work her husband, narrowing the childcare gap; the estimated household instead has her withdraw, widening the gender gaps in work and childcare. Because automation disproportionately depresses men's wages in the lower-middle of the distribution, it pushes a significant mass of couples into the region where the norm binds, so the same shock that narrows the gender wage gap widens the gender gaps in time use.

The findings point to two policy margins. Standard adjustment assistance retrains displaced workers; this paper shows that a fall in the husband's wage can trigger a secondary withdrawal of the wife, so wage insurance that stabilizes displaced men's earnings would also protect female labor supply by keeping households away from the region where the identity penalty binds, though the counterfactuals price it as the costlier route: a credit on the wife's own earnings buys the same hours at a fraction of the outlay. And because the shock operates through the gender composition of local labor demand, place-based policy that eases men's transition into service-sector work would attack the friction at its source instead of compensating for it.

Three margins remain open for future work. The model is static: a life-cycle version with human-capital accumulation would let the temporary withdrawal documented here feed back into the wife's later potential wages, and with them into the norm's long-run cost, along the lines of the dynamic collective models of Mazzocco (2007) and Voena (2015). The wedge is held fixed where the heterogeneity evidence says its salience varies across space, and automation also moves marriage formation and fertility (Anelli et al., 2024); modeling the norm's local intensity, together with the matching margin it operates on, would tie the size of the friction to the environments that sustain it. And couple-level diaries, which the ATUS does not collect, would allow the norm's incidence to be estimated household by household.

# References


Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Adão, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4):1949–2010.

Akerlof, G. A. and Kranton, R. E. (2000). Economics and identity. *The Quarterly Journal of Economics*, 115(3):715–753.

Andrews, I., Gentzkow, M., and Shapiro, J. M. (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, 132(4):1553–1592.

Anelli, M., Giuntella, O., and Stella, L. (2024). Robots, marriageable men, family, and fertility. *Journal of Human Resources*, 59(2):443–469.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–68.

Baker, M., Gruber, J., and Milligan, K. (2008). Universal child care, maternal labor supply, and family well-being. *Journal of Political Economy*, 116(4):709–745.

Bauernschuster, S. and Schlotter, M. (2015). Public child care and mothers' labor supply: Evidence from two quasi-experiments. *Journal of Public Economics*, 123:1–16.

Becker, G. S. (1965). A theory of the allocation of time. *The Economic Journal*, 75(299):493–517.

Bertrand, M., Kamenica, E., and Pan, J. (2015). Gender identity and relative income within households. *The Quarterly Journal of Economics*, 130(2):571–614.

Bick, A. and Fuchs-Schündeln, N. (2018). Taxation and labor supply of married couples across countries: A macroeconomic analysis. *American Economic Review*, 108(6):1543–76.

Binder, A. J. and Lam, D. (2022). Is there a male-breadwinner norm? The hazards of inferring preferences from marriage market outcomes. *Journal of Human Resources*, 57(6):1885–1914.

Blundell, R., Chiappori, P.-A., and Meghir, C. (2005). Collective labor supply with children. *Journal of Political Economy*, 113(6):1277–1306.

Blundell, R., Dias, M. C., Meghir, C., and Shaw, J. (2016). Female labor supply, human capital, and welfare reform. *Econometrica*, 84(5):1705–1753.

Blundell, R. and MaCurdy, T. (1999). Labor supply: A review of alternative approaches. In Ashenfelter, O. C. and Card, D., editors, *Handbook of Labor Economics*, volume 3A, chapter 27, pages 1559–1695. Elsevier.

Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2018). Children, Time Allocation, and Consumption Insurance. *Journal of Political Economy*, 126(S1):73–115.

Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1):181–213.

Borusyak, K., Hull, P., and Jaravel, X. (2025). A practical guide to shift-share instruments. *Journal of Economic Perspectives*, 39(1):181–204.

Bound, J., Brown, C., and Mathiowetz, N. (2001). Measurement error in survey data. In Heckman, J. J. and Leamer, E., editors, *Handbook of Econometrics*, volume 5, pages 3705–3843. Elsevier.

Brines, J. (1994). Economic dependency, gender, and the division of labor at home. *American Journal of Sociology*, 100(3):652–688.

Browning, M. and Chiappori, P.-A. (1998). Efficient intra-household allocations: A general characterization and empirical tests. *Econometrica*, 66(6):1241–1278.

Chetty, R., Guryan, J., Kroft, K., and Looney, A. (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive margin. *American Economic Review*, 101(3):471–75.

Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy*, 100(3):437–467.

Cortes, P., Feng, Y., Guida-Johnson, N., and Pan, J. (2024). Automation and gender: Implications for occupational segregation and the gender skill gap. NBER Working Paper 32030, National Bureau of Economic Research.

Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 19(6):3104–3153.

de Vries, G. J., Gentile, E., Miroudot, S., and Wacker, K. M. (2020). The rise of robots and the fall of routine jobs. *Labour Economics*, 66:101885.

Del Boca, D., Flinn, C., and Wiswall, M. (2014). Household Choices and Child Development. *The Review of Economic Studies*, 81(1):137–185.

Doepke, M. and Kindermann, F. (2019). Bargaining over babies: Theory, evidence, and policy implications. *American Economic Review*, 109(9):3264–3306.

Duffie, D. and Singleton, K. J. (1993). Simulated moments estimation of markov models of asset prices. *Econometrica*, 61(4):929–952.

Eissa, N. and Hoynes, H. W. (2004). Taxes and the labor market participation of married couples: The earned income tax credit. *Journal of Public Economics*, 88(9-10):1931–1958.

Foerster, H. (2025). Untying the knot: How child support and alimony affect couples' dynamic decisions and welfare. *Review of Economic Studies*, 92(5):3029–3066.

Frazis, H. and Stewart, J. (2004). What can time-use data tell us about hours of work? *Monthly Labor Review*, 127(12):3–9.

Friedberg, L. and Webb, A. (2006). Determinants and consequences of bargaining power in households. NBER Working Paper 12367, National Bureau of Economic Research.

Ge, S. and Zhou, Y. (2020). Robots, computers, and the gender wage gap. *Journal of Economic Behavior & Organization*, 178:194–222.

Giuntella, O., Lu, Y., and Wang, T. (2025). How do workers adjust to robots? Evidence from China. *The Economic Journal*, 135(666):637–652.

Glass, J. and Nath, L. E. (2006). Religious conservatism and women's market behavior following marriage and childbirth. *Journal of Marriage and Family*, 68(3):611–629.

Gobbi, P. E. (2018). Childcare and commitment within households. *Journal of Economic Theory*, 176:503–551.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.

Graetz, G. and Michaels, G. (2018). Robots at Work. *The Review of Economics and Statistics*, 100(5):753–768.

Grönqvist, E., Hensvik, L., Okuyama, Y., and Thoresson, A. (2026). Women's labor market opportunities and equality in the household. *Journal of Public Economics*, 254:105557.

Guiso, L., Sapienza, P., and Zingales, L. (2003). People's opium? Religion and economic attitudes. *Journal of Monetary Economics*, 50(1):225–282.

Havnes, T. and Mogstad, M. (2011). Money for nothing? Universal child care and maternal employment. *Journal of Public Economics*, 95(11-12):1455–1465.

Hederos, K. and Stenberg, A. (2022). Gender identity and relative income within households: Evidence from Sweden. *Scandinavian Journal of Economics*, 124(3):744–772.

Heggeness, M. and Murray-Close, M. (2019). Manning up and womaning down: How husbands and wives report earnings when she earns more. Opportunity and Inclusive Growth Institute Working Papers 28, Federal Reserve Bank of Minneapolis.

Ichino, A., Olsson, M., Petrongolo, B., and Skogman Thoursie, P. (2026). Taxes, childcare, and gender identity norms. *Journal of Labor Economics*, 44(2):553–586.

Juster, F. T. and Stafford, F. P. (1991). The allocation of time: Empirical findings, behavioral models, and problems of measurement. *Journal of Economic Literature*, 29(2):471–522.

Kan, M. Y. and Pudney, S. (2008). Measurement error in stylized and diary data on time use. *Sociological Methodology*, 38(1):101–132.

Keane, M. P. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature*, 49(4):961–1075.

Kleibergen, F. and Paap, R. (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics*, 133(1):97–126.

Kletzer, L. G. and Litan, R. E. (2001). A prescription to relieve worker anxiety. Policy Brief 01-2, Institute for International Economics.

Knowles, J. (2013). Why are married men working so much? An aggregate analysis of intra-household bargaining and labor supply. *The Review of Economic Studies*, 80(3):1055–1085.

Lippmann, Q., Georgieff, A., and Senik, C. (2020). Undoing gender with institutions: Lessons from the German division and reunification. *The Economic Journal*, 130(629):1445–1470.

Mazzocco, M. (2007). Stochastic dynamic models of the household. *Journal of Political Economy*, 115(6):890–935.

Robinson, J. P. and Bostrom, A. (1994). The overestimated workweek? What time diary measures suggest. *Monthly Labor Review*, 117(8):11–23.

Rosen, J. B. (1965). Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534.

Ruggles, S., Flood, S., Sobek, M., Backman, D., Chen, A., Cooper, G., Richards, S., Rodgers, R., and Schouweiler, M. (2024). IPUMS USA: Version 15.0 [dataset]. Minneapolis, MN: IPUMS.

Stephens Jr., M. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics*, 20(3):504–537.

Voena, A. (2015). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review*, 105(8):2295–2332.

Webb, M. (2019). The Impact of Artificial Intelligence on the Labor Market. Working paper, Stanford University. SSRN 3482150.

Zinovyeva, N. and Tverdostup, M. (2021). Gender identity, coworking spouses, and relative income within households. *American Economic Journal: Applied Economics*, 13(4):258–84.

**Funding** This project has been funded by the Belgian National Fund for Scientific Research (FNRS).

**Data availability** The following data sources have been used for the analysis in this article:

- **American Time Use Survey (ATUS)** – U.S. Bureau of Labor Statistics, 2003–2023.
- **Current Population Survey (CPS)** – ASEC and basic monthly files, via IPUMS-CPS (Flood et al. 2024 release); 1990–2023.
- **1990 U.S. Census Microdata** – Used for baseline potential hourly wages, via IPUMS-USA.
- **Churches and Church Membership in the United States, 1990** – Association of Statisticians of American Religious Bodies (ASARB).
- **ISSP "Family and Changing Gender Roles IV"** – 2012 U.S. subsample.
- O*NET manual-task intensity scores – Autor & Dorn (2013, *AER* 103(5): 1553-97).
- 1990 commuting-zone × industry employment shares – Acemoglu & Restrepo (2020, *JPE* 128(6): 2188-2244).
- 1970 commuting-zone × industry employment shares – Decennial IPUMS data.
- Industrial-robot stocks (country × industry, 1993-2019) – International Federation of Robotics, *World Robotics* database.
- 1990 industry employment for EU-5 (used in the robot instrument) – EU KLEMS Growth & Productivity Accounts, 2017 edition.
- U.S. import-exposure, domestic absorption and baseline instrument files – Autor, Dorn & Hanson (2013, "The China Syndrome," *AER* 103(6): 2121-68).
- Chinese exports to eight alternative OECD markets – UN Comtrade extraction following Autor–Dorn–Hanson.
- Commuting-zone boundaries and county cross-walks – Autor, Dorn & Hanson replication materials (2013).

**Conflict of interest** The author declares that he has no conflict of interest.

# Appendix

## A Discussion of the modeling choices

The model is designed around two regularities of the data: the prevalence of corner solutions in daily childcare, where one parent reports no quality time, and the reversal of the time-use profiles at the equal-earnings threshold (Section 2.1). Each modeling choice serves one of them.

**Non-cooperative childcare.** Treating market hours as contractible but childcare as a privately provided public good generates corners without ad hoc restrictions. With logarithmic utility and Cobb-Douglas child quality, the cross-partial derivative of the spouses' childcare times is zero, so each parent's care depends only on his or her own shadow price of time, which is set by first-stage market hours; a parent provides no quality time whenever that shadow price exceeds the private marginal benefit of care. The semi-cooperative equilibrium under-provides care relative to the Pareto-efficient benchmark, in which each spouse internalizes the partner's valuation of child quality. That benchmark is not assumed away: it is estimated as a competitor in Section 4.4, where even a bargaining weight acting directly on care and free to kink at parity fails to reproduce the paternal-care inverse-V (Table 10, Proposition 1).

**The norm as a budget wedge.** Relative-status concerns conventionally enter as additive utility costs, in the tradition of Akerlof and Kranton (2000) and as operationalized by Bertrand et al. (2015). Two reasons favor a wedge on the budget constraint. The first is interpretive: $\phi(s)$ is the resource cost of maintaining a traditional breadwinner hierarchy when the wife's market value approaches or surpasses the husband's, in the form of the compensatory behaviors and domestic frictions that the "gender display" literature documents (Brines, 1994; Bertrand et al., 2015); the household forfeits consumable income to status preservation. The second is substantive: an additively separable identity cost shifts the *level* of utility but leaves the consumption–leisure marginal rate of substitution unchanged, so on its own it does not bend the optimal-hours schedule at parity. The budget wedge does, because it scales the marginal return to the wife's labor: as $s$ rises, the declining $\phi(s)$ acts as an endogenous marginal tax that depresses the net return to her hours enough to override the added-worker motive, so that she works less even though the household is poorer. This is the feature that reproduces the non-monotonic profiles of Section 2.1.

**A fixed bargaining weight.** Endogenizing the Pareto weight as an increasing function of the wife's wage share, $\theta(s)$, generates the monotonic rise in female leisure observed in the data, but continuous shifts in bargaining power cannot produce localized reversals in market work and childcare at the parity point. Holding $\theta$ fixed across the distribution and introducing the wedge isolates norm maintenance as the driver of the reversals; Section 4.4 lets distribution-factor weights of increasing flexibility compete with the wedge on the same moments.

**Potential wages.** Indexing the wedge to *potential* wages is both a modeling convenience and a substantive choice. It keeps $\phi(s)$ a constant in the optimization, preserving the closed-form solution and the uniqueness of equilibrium. Substantively, it means the norm acts on the wife's market *value*, a quantity she cannot manipulate; withdrawing from work therefore does not relax the penalty. This is the more demanding "gender display" interpretation (Brines, 1994; Bertrand et al., 2015) and it makes the wife's withdrawal a pure utility loss. A specification in which $\phi$ instead responded to the realized labor-income share would add a mechanical feedback (withdrawing lowers that share and so loosens the penalty) that would only *amplify* the contraction in female market work. The potential-wage wedge is therefore the conservative choice for bounding the norm's effect.

**A continuous wedge.** The functional form $\exp[-\chi s^{\kappa}]$ lets the data determine the onset and the intensity of the friction, with no discrete penalty imposed at an arbitrary cutoff, keeps consumption strictly positive, and preserves the differentiability that delivers a well-defined marginal rate of substitution and a unique equilibrium in market hours.

**From assumptions to patterns.** Three mappings summarize how the model produces the profiles it is estimated on. (i) Logarithmic utility with Cobb-Douglas child quality produces the corner solutions: parents stop providing quality time when market hours raise the shadow price of their time above the private marginal benefit of care. (ii) The wedge on the wife's effective wage produces the kink in her market hours: below the onset the wage share moves her hours through comparative advantage and income, past it the falling marginal return $\phi(s)w_f$ dominates. (iii) The second-stage responses to first-stage hours produce the childcare reversals: when the wife's hours fall her shadow price of time falls and her care rises, while the husband's hours rise and his care falls, the mirror-image pattern in the data that no cooperative weight can generate (Proposition 1).