Workplace Flexibility and the Motherhood Penalty: Evidence from the Diffusion of Remote Work*

Gaetano Basso, Maria De Paola, Salvatore Lattanzio, Matteo Paradisi

March 2026

Abstract

We study whether workplace flexibility is a key driver of the child penalty. We exploit the sharp and heterogeneous diffusion of work-from-home (WFH) contracts after COVID-19 in a difference-in-differences framework. Our analysis combines administrative data on the universe of Italian WFH contracts with matched employer–employee records, administrative mother–father links, fertility records, and firm-level balance sheets. Greater exposure to flexible work arrangements substantially reduces mothers' post-childbirth earnings losses, primarily through higher weeks worked, lower part-time incidence, lower take-up of parental leave, and improved career progression. The gains are larger among younger, lower-earning and commuting mothers, consistent with flexibility relaxing time and contractual constraints where they are most binding. IV estimates indicate that holding a WFH contract offsets a large share of the child penalty. While fathers' own earnings do not respond to flexibility around childbirth, fathers' exposure to flexible work reduces mothers' earnings losses by a comparable magnitude. This pattern points to household-level time constraints—rather than mother-specific constraints—as a central mechanism for the child penalty, and suggests a more limited role for explanations based primarily on employer discrimination or gender norms. Consistent with this interpretation, fertility—a joint household decision—increases for mothers more exposed to flexibility, because flexibility lowers the labor market cost of the marginal child. A counterfactual exercise shows that the life-cycle widening of the gender earnings gap would have been 13% smaller under current WFH diffusion and up to 26% smaller if all remotable jobs had adopted flexible arrangements.

Keywords: Work-from-home; Child penalty; Gender earnings gap; Fertility; Labor market adjustment

***JEL* codes:** J13, J16, J22, J31

*Contact information: Gaetano Basso, Banca d'Italia, gaetano.basso@bancaditalia.it; Maria De Paola, University of Calabria, INPS and IZA, m.depaola@unical.it; Salvatore Lattanzio, Banca d'Italia, RFBerlin and Dondena, salvatore.lattanzio@bancaditalia.it; Matteo Paradisi, EIEF and CEPR, matteo.paradisi@eief.it. The views expressed in this paper are those of the authors only and should not be attributed to the institutions with which they are affiliated.

1 Introduction

Despite major progress in women's education and labor market participation, substantial gender gaps in employment and earnings persist across advanced economies. A large body of evidence attributes much of these gaps to the arrival of children, which triggers a persistent divergence in men's and women's career trajectories—the so-called *child penalty* (Angelov et al., 2016; Casarico and Lattanzio, 2023; Kleven et al., 2019). These penalties partly arise from the difficulty of reconciling work and care within organizational structures that rely on physical presence and rigid schedules (Goldin, 2014). Yet it remains unclear whether the child penalty primarily reflects such contractual rigidities, or instead stems from deeper forces such as gender norms, preferences over specialization, or employer discrimination. In this paper, we investigate whether work flexibility is a central driver of the child penalty and quantify the extent to which the recent expansion of work-from-home (WFH) arrangements has mitigated it.

The rapid diffusion of WFH during the COVID-19 pandemic provides a unique opportunity to assess this question. By altering *where* and *when* work can be performed, remote arrangements relax constraints embedded in traditional workplace organization. WFH reduces commuting time (Le Barbanchon et al., 2020) and allows tasks to be completed outside the physical workplace (Barrero et al., 2023; Basso et al., 2025; Harrington and Kahn, 2025), potentially enabling parents—especially mothers—to remain more strongly attached to their jobs around childbirth. At the same time, the implications of this shift are theoretically ambiguous: remote work may also exacerbate gender asymmetries in unpaid care or limit career advancement if physical presence remains valued (Cullen and Perez-Truglia, 2023). Whether WFH ultimately narrows or reinforces the child penalty is therefore an empirical question that speaks directly to the underlying sources of gender inequality after childbirth.

This paper examines whether and to what extent access to WFH mitigates the child penalty by exploiting variation in exposure to remote work across sectors and locations. Beyond estimating the impact of a specific work arrangement, we use the diffusion of WFH to assess whether rigid workplace organization is a central driver of the earnings divergence that follows childbirth. In doing so, the paper connects two strands of literature—on the sources of the child penalty and on the labor market consequences of remote work—and provides new evidence on the role of contractual flexibility in shaping gender inequality.

Our analysis combines administrative data on the universe of Italian WFH contracts with matched employer–employee social security records and administrative links between mothers and fathers. To our knowledge, this is the first paper to merge all these sources. The richness of the data allows us to observe the exact timing of childbirth, track mothers' and fathers' labor market outcomes before and after childbirth, examine

cross-parent spillovers, future fertility, and measure how contractual arrangements and firm-level outcomes evolve over time.

Our empirical strategy builds on recent studies that exploit the sharp post-COVID diffusion of WFH (e.g., Barrero et al., 2023; Basso et al., 2025; Lee, 2023). We estimate a difference-in-differences model on repeated cross-sections of mothers to identify how their labor outcomes change in sector-locations that were more exposed to remotable jobs once flexible arrangements became widespread. We measure this exposure using an index of remotability that captures the pre-COVID share of jobs in each sector that could plausibly be performed remotely, which we interact with pre-pandemic internet broadband coverage to proxy for local exposure to WFH. Our identification relies on parallel trends between more and less exposed sector-locations. We provide supportive evidence for this assumption by looking at the pre-pandemic trends in mothers' labor outcomes across differently exposed sector-locations. Moreover, because our main focus is the medium-run response of mothers' outcomes rather than the short-term one, and because we find no effect of exposure on firm-level performance, it is unlikely that contemporaneous COVID-related sector-location shocks confound our results. A potential concern in our strategy that relies on repeated cross-sections is, however, that the composition of mothers may change over time as the availability of WFH affects fertility and career choices. To address this, we use pre-birth earnings as a proxy for a mother's earnings potential, allowing us to net out selection effects and isolate the true reduction in the child penalty attributable to flexible work.

We begin by documenting the reduced-form effects of exposure to remotable work on mothers' labor market outcomes. Moving from the 10^{th} to the 90^{th} percentile of exposure (equivalent to comparing a web-related job to a manufacturing one) raises mothers' earnings in the year after childbirth by about 980 euros. This gain is accompanied by an increase of 0.53 weeks worked and a higher probability of full-year employment, together with a decline in part-time work and in parental leaves. These patterns indicate that greater flexibility expands mothers' feasible labor supply set after childbirth, allowing them to maintain stronger attachment along both intensive and extensive margins. We do find some evidence that mothers in highly exposed sector-locations have higher pre-birth earnings and weeks worked after 2020, consistent with mild changes in their selection. However, when we net out these pre-birth gains, the increase in post-childbirth outcomes remains positive and significant, implying that flexible arrangements attenuate the earnings loss associated with childbirth rather than merely reshuffling who becomes a mother.

Next, we translate the exposure estimates into the causal effect of holding a WFH contract by instrumenting the probability of having such a contract with sector-location exposure. The first stage is strong: moving from the 10^{th} to the 90^{th} percentile of exposure increases the likelihood of having a WFH contract by 13 percentage points. Holding a

WFH contract reduces the average earnings drop around childbirth by 87% of the average drop or by 27% of the average earnings before childbirth. Using pre-birth earnings as a measure of potential earnings, we compute the implied reduction in the child penalty that accounts for changes in the selection of mothers. Depending on the calibration of the baseline penalty, WFH contracts reduce the percentage drop in earnings by between 31 and 36 percentage points—equivalent to a reduction of the short-run penalty between 43% and 72%. These magnitudes suggest that rigid work arrangements account for a substantial share of the earnings divergence following childbirth.

To support the causal interpretation of the IV estimates, we examine possible violations of the exclusion restriction. We show that sector-location exposure has no significant effect on changes around childbirth in mothers' employers' outcomes such as size, sales, and value added (both in levels and per worker). Similarly, we find no effect on changes in coworkers' or male coworkers' earnings, indicating that the results are not driven by adjustments in rent-sharing. Mothers' outcomes thus improve relative to their coworkers', consistent with the mechanism operating through their contractual work arrangements rather than through broader firm-level shocks.

We then study heterogeneity in the effects across subgroups of mothers. Flexible arrangements generate the largest gains for mothers facing the most binding time constraints around childbirth. Younger mothers benefit disproportionately: relative to pre-birth earnings, the difference between mothers below 30 and those aged 36 and above amounts to 42 percentage points in the IV. Lower-earning mothers also gain more, both in absolute terms and relative to pre-birth earnings. This is not driven by a stronger first stage, indicating that a given increase in flexibility translates into larger labor market gains when baseline earnings are lower. A similar pattern holds by father's earnings: the IV effect is larger when the partner is a lower earner, consistent with tighter household-level constraints at the bottom of the income distribution. Finally, the effects are larger for mothers who commuted to work before childbirth, a group for whom flexible arrangements most directly free up time. Taken together, these results confirm that rigid work arrangements are a key driver of the child penalty: flexibility generates the largest gains precisely where the labor supply costs of workplace rigidity appear most severe, compressing inequality among mothers and mitigating the post-childbirth rise in earnings dispersion.

We also examine fathers' labor market responses to increased exposure to flexible work arrangements. Consistent with the well-established absence of a child penalty for fathers, we find no significant effect of remotability on fathers' earnings around childbirth. Fathers' labor supply remains largely unchanged. However, fathers' exposure to flexible work has a substantial impact on mothers' outcomes: when fathers work in more remotable sectors, mothers' post-childbirth earnings improve by an amount comparable to the effect of mothers' own exposure. The absence of fathers' earnings responses, combined

with the sizable effect on mothers, suggests that flexibility operates primarily through intra-household time reallocation rather than through increases in fathers' labor supply. These findings indicate that the child penalty is not solely a mother-specific labor market phenomenon but instead reflects a joint household time constraint interacting with rigid workplace arrangements for both spouses.

Adding to the evidence that workplace rigidity operates as a household-level constraint, we also examine whether the diffusion of flexible work arrangements affected fertility decisions. If the child penalty reflects the labor market cost of motherhood, greater flexibility lowers the marginal cost of an additional child and may therefore raise fertility. To test this, we follow over four years two cohorts of working women—one observed before and one after the COVID shock—and compare their fertility outcomes across sector-locations with different remotability. We find a significant increase in births after the expansion of WFH, both on the intensive margin (*i.e.* higher-order births among women who already had at least one child) and on the extensive margin of entry into motherhood. The effect is smaller on the latter, consistent with flexible arrangements only partially offsetting the fixed costs of a first child. We then aggregate to the sector–location level and find that the total number of working mothers rises significantly in more remotable cells, driven by women aged 27 and above. These results indicate that rigid work arrangements were distorting not only mothers' post-birth earnings trajectories but also household family formation decisions, reinforcing the interpretation that workplace rigidity is a key driver of the child penalty.

Finally, we assess the implications of the diffusion of flexible work arrangements for the gender earnings gap over the lifecycle. Focusing on cohorts born between 1974 and 1979—who reached the end of their fertile years before COVID-19—we study the within-cohort evolution of the female-to-male earnings ratio. In the data, this ratio falls by 20 percentage points between ages 25 and 40. When we remove the full child penalty, the ratio remains flat or even increases slightly, confirming that the penalty explains nearly the entire widening of the gap. Under the current post-COVID WFH adoption, the drop in the ratio would shrink by 2.4 percentage points, equivalent to 13% of the observed lifecycle change. If all remotable jobs offered WFH contracts, the gap would narrow by a total of 5.0 percentage points (26% of the observed drop). Although WFH cannot close the entire gender earnings gap, it can mitigate its growth over the lifecycle. In complementary analysis, we show that mothers in our focus cohorts were not disproportionately concentrated in remotable sectors, while the post-COVID fertility increased precisely in those sectors. This suggests that future cohorts, whose fertility and career choices align more closely with remote work opportunities, may experience larger gains.

We contribute to two strands of the literature. First, we contribute to the literature on the determinants of the child penalty (Angelov et al., 2016; Kleven et al., 2024a,b, 2019), which has documented large and persistent gender gaps arising at childbirth but

has provided more limited causal evidence on their underlying drivers. A prominent explanation emphasizes the role of workplace organization and the lack of temporal flexibility in jobs (Goldin, 2014, 2022), alongside institutional and household factors shaping labor supply and specialization (Cortés and Pan, 2023; Olivetti and Petrongolo, 2016, 2017), and more broadly reflects a central theme in the recent literature on gender inequality: the distinction between preferences and constraints in shaping gender gaps. In this perspective, much of the remaining inequality in advanced economies is understood to arise from differential work–family trade-offs and constraints associated with childcare, rather than intrinsic differences in preferences or skills (Olivetti et al., 2024). Workplace amenities—such as flexibility and remote work—are therefore key, as they relax these constraints and influence job sorting, labor supply, and earnings dynamics. However, most existing evidence relies on survey data or on education- and occupation-based proxies for job flexibility, making it difficult to isolate the causal role of contractual work arrangements. Pre-pandemic US studies show that greater feasibility of working from home reduces the motherhood employment penalty but may also induce sorting into lower-paying firms (Harrington and Kahn, 2025; Jack et al., 2025). In Europe and Latin America, measures of remote-work feasibility have been associated with narrower gender gaps in hours and employment (Arntz et al., 2022; Crescenzi et al., 2025; Zarate, 2025), while pandemic-era evidence is more mixed, likely reflecting contemporaneous childcare disruptions and shifts in unpaid work (Dunatchik et al., 2021; Farooqi, 2023; Goldin, 2022; Heggeness and Suri, 2021; Pabilonia and Vernon, 2022; Song, 2025). In contrast, we combine matched employer–employee administrative records with detailed data on the actual adoption of work-from-home contracts and exploit their rapid post-COVID diffusion as a source of plausibly exogenous variation in workplace flexibility. The richness of the data allows us to trace the effects of remote work not only on mothers' earnings, employment, and career progression, but also at the household level. We observe both parents before and after childbirth, examine cross-parent spillovers, and study how flexibility affects selection into motherhood, intra-household adjustments, and fertility decisions. This comprehensive perspective enables us to quantify how workplace flexibility reshapes the labor market cost of children within households and to assess its implications for the life-cycle evolution of the gender earnings gap. Most importantly, by isolating the causal effect of a specific contractual feature of jobs, we provide direct evidence that rigid workplace organization is a central driver of the child penalty.

Second, we contribute to the literature on WFH and worker outcomes by shifting attention to the margins that generate persistent gender inequality after childbirth. Existing studies—often based on experimental settings or evidence from single firms—show that remote and hybrid work affect retention, satisfaction, and career progression, while a growing body of work documents their rapid diffusion in the post-pandemic period (Atkin et al., 2023; Barrero et al., 2023; Bloom et al., 2024, 2015; Choudhury et al., 2024; Mas and

Pallais, 2017). Recent work using employer–employee data examines selection into remote work and its effects on earnings, promotions, productivity, and health (Emanuel and Harrington, 2024; Fenizia and Kirchmaier, 2025; Gibbs et al., 2023; Goux and Maurin, 2025). Building on this literature, we study WFH diffusion as an economy-wide shock to flexibility and show that its effects are highly state-dependent, concentrated around parenthood rather than reflecting a uniform remote-work premium or penalty. Moreover, the extent to which remote work mitigates the mother–child penalty depends on household constraints, including the degree of contractual flexibility available to fathers. Our analysis is closely related to Aksoy et al. (2026), who, using survey data, document a positive association between working from home and realized, planned, and lifetime fertility. While their approach provides valuable descriptive evidence based on cross-sectional variation in remote work arrangements, we complement this perspective by exploiting the sudden, economy-wide expansion of formal WFH contracts and implementing an individual-level difference-in-differences design using population-wide administrative data. This allows us to provide causal evidence on the relationship between remote work and fertility outcomes. Moreover, by studying fertility jointly with post-birth labor supply and earnings trajectories within couples, we offer a more comprehensive assessment of how rigid work arrangements shape household behavior and contribute to the mother–child penalty.

2 Conceptual Framework

This section develops the economic framework underlying our empirical analysis. It clarifies the mechanisms through which workplace flexibility affects the child penalty and provides a guide to interpreting the empirical results. A formal model is presented in Appendix A.

Flexibility and the Child Penalty. Childbirth raises both the time devoted to childcare and the coordination costs of combining work and care. Rigid workplace arrangements—fixed schedules, commuting requirements, low adaptability to childcare logistics—tighten a time feasibility constraint, generating a child penalty in earnings. Greater remotability reduces both commuting and scheduling frictions, mitigating these losses. This mechanism yields three testable predictions. First, if these frictions are relevant, greater exposure to flexible work should translate into higher mothers' post-birth earnings—so the earnings response to flexibility serves as an indirect test of whether rigid arrangements were a binding constraint. Second, the adjustment should operate primarily through higher weeks worked, lower part-time incidence, and stronger labor market attachment. Third, the effect should appear as an immediate level shift around childbirth rather than only through slower changes in earnings growth and job continuity.

Heterogeneity as a Diagnostic Tool. Heterogeneity in the flexibility effect could serve as a diagnostic for the sources of the child penalty. If the penalty is driven by rigid workplace arrangements, it should bite harder where rigidity is more severe—in lower-paying jobs with less schedule autonomy and greater coordination burdens. Finding larger *relative* gains from flexibility for lower-earning mothers would therefore be consistent with job rigidity as a mechanism. Finding larger *absolute* gains would be even stronger evidence, as it would require that lower-earning mothers exhibit larger labor supply responses—for instance because rigid arrangements force them closer to discrete margins such as part-time work or non-employment.

Age heterogeneity speaks instead to the role of career disruption. If the penalty operates partly through interruptions to early-career accumulation, where wage growth is steeper and labor market attachment is more fragile, then flexibility should matter more for younger mothers, both by relaxing contemporaneous work–care frictions and by preserving continuity at the stage where it is most valuable. Finding a stronger flexibility effect for younger mothers would therefore point to career disruption, rather than only a static hours adjustment, as a source of the child penalty.

Why the Household Perspective Matters. Extending the framework to a household setting reframes childbirth as a household-level time constraint rather than a purely individual shock. Three forces explain why this constraint nonetheless falls disproportionately on mothers: (*i*) comparative advantage may push childcare toward the lower-wage parent; (*ii*) some early childcare needs—recovery, breastfeeding, mandatory leave—are difficult to reallocate, and asymmetric initial disruptions can persist if there are returns to employment continuity; and (*iii*) gender norms may reinforce both. These channels all rationalize a mother-side child penalty, but carry different implications for the role of fathers' workplace flexibility. This makes fathers' exposure an informative test. If the penalty were driven mainly by mother-side frictions—employer discrimination, stigma, or other mechanisms operating independently of household time constraints—then fathers' flexibility would have no bearing on mothers' earnings. If instead the relevant friction is also a household-level time constraint interacting with rigid workplace arrangements, greater flexibility on the father side should relax that same constraint by facilitating intra-household childcare reallocation.

Fathers, Zero Child Penalty, and Intra-household Reallocation. The empirical absence of a father child penalty is itself informative. It suggests that employer-side rigidities, such as full-time norms, hours constraints, and promotion premia, limit fathers' ability to reduce market hours around childbirth. Consistent with this interpretation, evidence showing no effect of working from home on fathers' earnings indicates that even when flexibility increases, fathers have little scope to adjust labor supply or career

trajectories. At the same time, father-side flexibility is associated with improvements in mothers' outcomes, pointing to within-household reallocation of time and responsibilities. Taken together, the coexistence of a null father penalty, no earnings response to father flexibility, and a positive effect of father flexibility on mothers' outcomes suggests that the child penalty arises from a household-level time constraint interacting with rigid labor contracts. This pattern is difficult to reconcile with mechanisms that operate exclusively on the mother side.

Fertility. Flexibility should affect not only labor supply but also fertility. If rigid work arrangements raise the expected cost of an additional child by tightening time and coordination constraints, greater flexibility reduces the marginal cost of childbirth. A positive fertility response among more-exposed households would reinforce the household-constraint interpretation, suggesting that pre-COVID rigidities were distorting family choices, not only post-birth earnings. This margin is especially informative at higher parity, where the decision depends more directly on whether existing arrangements can absorb additional childcare needs rather than on anticipatory career concerns.

3 Data

3.1 Administrative Data on Workers and Firms

Universe of Social Security Data. We use matched employer-employee administrative data covering the years 2014 to 2024 from the Italian Social Security Archives (INPS). The dataset includes the universe of private sector employees. [1] Our analysis focuses on childbirth episodes happening to the universe of employees. The data provide rich information on employment relationships, including annual labor earnings, number of weeks worked (both raw and in full-time equivalent terms), type of contract (permanent, fixed-term, or seasonal), work schedule (full-time or part-time), municipality of work, and a broad occupation classification (blue-collar, white-collar, manager, apprentice or other). We observe the date of entry into the firm, as well as the end date and reason for termination, when applicable.

The dataset also includes demographic information on workers—gender, year of birth, migrant status, and province of residence—as well as detailed firm-level information. For all firms, we observe the sector of activity and the dates of establishment and (if applicable) closure. For incorporated firms, the dataset is enriched with balance sheet information from Orbis, including value added and revenues.

[1] We are not able to observe job spells in the public sector and self-employed work.

Chilbirth Episodes and Household Linkages. To identify childbirth episodes, we combine two data sources. The primary source is administrative records from the *Assegno Unico Universale* (Universal Childcare Allowance, UCA), a universal child subsidy available to all families with children under age 18. For each payment, we observe the social security numbers of the mother, the father, and the child, as well as the child's date of birth. We use the latter to assign a childbirth year to each mother in the matched employer–employee records. Because the data include identifiers for both parents, they allow us to link mothers and fathers and to match couples within the employer–employee data.

The take-up rate is about 95 percent for children born between 2015 and 2023. We therefore recover paternal identifiers and the child's year of birth for the vast majority of maternity events in our sample, making the UCA records a reliable proxy for births in Italy over this period. Since applying involves a modest administrative burden and the benefit is means-tested, the linked sample slightly overrepresents lower-income households. To address this concern, we complement the UCA information with data on the universe of maternity leaves taken during the same period. Maternity leave is mandatory for employed mothers at the time of childbirth, allowing us to identify the full set of maternity episodes among working mothers—our main population of interest. These records do not contain information on fathers, so analyses involving fathers rely on the UCA-based sample. Reassuringly, our main results on mothers' outcomes are qualitatively and quantitatively similar when restricted to this subsample, indicating that the slight overrepresentation of lower-income households in the UCA data does not drive our findings.

Sample of Workers' Careers. Finally, we also use a sample of workers sourced from the same Social Security matched employer-employee archives, and available at Banca d'Italia for the period 1990-2022. These data contain a random sample of workers born in 24 dates of each year. The data contain the same information as the one available for the universe of matched employer-employee records. The advantage of these data is that it can be linked to the historical contribution records of individuals (even before 1990), which allow to reconstruct the full life cycle of events and therefore provides information on maternity episodes happening before our observation window. We use these data—whose sample size is of 125,327 mothers—to study the implications of work from home on the lifecycle evolution of the gender earnings gap, as explained in Section 8.

3.2 Work From Home Arrangements in the Administrative Data

We access data on the universe of work from home contracts in Italy. This information comes from the "lavoro agile" declarations that have been introduced by the Italian gov-

ernment in 2017. "Lavoro agile" is defined by the Law a form of subordinate employment characterized by the absence of time and space constraints; it thus corresponds to WFH. This type of contracts was introduced with the specific goal of helping workers balance work and family life while simultaneously fostering increased productivity. Mandatory reporting of WFH spells at the worker and week level are required by Law as workers are entitled to public insurance for accidents and occupational diseases, even in relation to services rendered outside the company premises. They are consistently recorded by firms through an online portal since 2020.

For each WFH spell, we observe the firm identifier, the gender, year of birth and citizenship of the employee, and the week the spell starts and ends. Thanks to this information, we are able to match these data to social security records. We aggregate the data at the year level and create firm-level cells based on gender-age-citizenship combinations and we use these cells as the unit for the match. When there are more employees in social security records than WFH contracts in a given cell, we attach a probability of holding such contract to the employees in this group. We compute this probability dividing the number of WFH contracts in the cell by the total number of employees within it. In 15% of cases, we find one-to-one matches between WFH data and social security records. In 77% of cases, we have more employees in the cell than the number of WFH contracts, but in 11% of these cases the discrepancy is of at most two employees. For each individual, we therefore define a probability of holding a WFH contract that we will use as an endogenous variable in our analysis. In the 8% of cases in which the cell is smaller than the number of WFH contracts, we winsorize the probability to one.

3.3 Baseline Sample Construction

We focus on the sample of mothers who i) gave birth for the first time between 2016 and 2023, ii) were between 18 and 45 years old at childbirth, and iii) were employed in the private sector at the time of childbirth. For each mother, we only consider the first childbirth. Table 1 reports descriptive statistics in Column 1 for the 796,174 mothers who satisfy these restrictions. Column 2 reports descriptives for the full sample of 938,667 fathers observed in UCA data. Columns 3 and 4 then report the corresponding statistics for the sample of linked households, which includes 428,397 mothers and fathers both employed in the private sector at the time of childbirth.

Relative to fathers, mothers in our baseline sample are younger on average (32.5 vs. 34.3 years), less likely to work full-time (59 vs. 83 percent), and earn less (16,425€ vs. 19,981€). They are also less likely to hold managerial positions and are more concentrated in smaller firms, while fathers are more concentrated in industry and in work environments with higher pay and higher sales and value-added per worker. These differences look

similar when we move to the linked-household sample, suggesting that the household-level analysis is conducted on a subset that remains broadly comparable in its main observable features.

4 Empirical Strategy

Our empirical strategy exploits the sudden expansion of WFH contracts at the onset of the Covid-19 pandemic. We first construct a measure of heterogeneous exposure to this shock across sector–location combinations, then describe how we use it to instrument WFH adoption in our empirical specifications, and discuss the main identification assumptions and threats.

Exposure to Remote Work. WFH in Italy was negligible before the pandemic—in 2019, only 1.5% of non-farm private sector employees worked from home at least once a week on average—and expanded sharply after it, reaching 11.7% in 2021 before stabilizing around 8.5% (Figure B.1). The increase was more pronounced for women than for men.

To instrument the adoption of WFH contracts at the onset of the pandemic, we construct a sector–location index of exposure to remote work that interacts two predetermined components. The first is an index of WFH potential at the 4-digit sector level, defined as the pre-pandemic share of occupations feasible to perform remotely based on task content (Basso et al., 2022; Dingel and Neiman, 2020). This measure exhibits substantial heterogeneity across sectors: construction and retail display low potential, while information and communication, finance, and scientific services show much higher feasibility, consistent with actual WFH use in 2024 (Table B.1). The second is the quality of broadband infrastructure in the commuting zone (CZ) where the firm is located, measured as average download speed in 2019 from AGCOM data.[2]

We refer to the interaction of these two components as "Remotable". It captures variation in firms' ability to adopt remote work, driven by factors—task composition and local infrastructure—that were predetermined and outside firms' control when the pandemic hit.[3] Importantly, identification will not rely on cross-sectional differences in sectors or locations per se: all specifications will include sector-by-location fixed effects and time-varying two-digit sectoral trends to absorb differential dynamics unrelated to

[2]Beyond firm-level ICT investment, effective remote work requires high-speed broadband access at home and at the workplace. In 2019, broadband access varied substantially across CZs, with a 90th-to-10th percentile ratio of average speed of 10.3 and a mean of 72.4 Mbps. See Basso et al. (2025) and Boeri et al. (2025) for evidence on the role of ICT investment and Lamorgese et al. (2024) on the role of managerial practices in determining WFH adoption.

[3]Personal conditions correlated with labor supply decisions, including housing arrangement and household composition, also shape individual WFH take-up (Burdett et al., 2024). By interacting remotable task feasibility with local infrastructure, our measure aims to isolate the demand-side component of WFH adoption.

remotability. Under these controls, "Remotable" isolates plausibly exogenous variation in the ease of transitioning to remote work post-pandemic across sector–location cells.

Empirical Specification. To quantify the effects of interest, we run two specifications, both implementing a difference-in-differences that exploits two sources of variation: differential exposure to remote work across sector–location cells, and the sharp increase in WFH adoption following the COVID-19 pandemic. We run our analysis on the sample of mothers described in Section 3.3.

The first specification is a reduced-form equation that studies how being in a remotable job—*i.e.*, in a sector with a high WFH potential and in a CZ with high-speed internet—affects the mother outcomes once WFH contracts become available. The identifying variation comes from the fact that sector-locations with remotable tasks will be more prone to adopt these contracts when they become widely diffused after the pandemic shock. We estimate

$$Y_{it} = \alpha_{s(i)} + \lambda_t + \sum_{k=2016}^{2024} \beta_k \cdot I\left(t = k\right) \cdot \text{Remotable}_{s(i)} + \sum_{k=2016}^{2024} \gamma_k \cdot I\left(t = k\right) \cdot X_i + \varepsilon_{it} \quad (1)$$

where i indexes a mother, t is the calendar time at which she gives birth, α_s is a location×sector fixed-effect, λ_t is a calendar year fixed-effect, $\text{Remotable}_{s(i)}$ is the remotability of the job held by i as defined above, and X_i is a set of individual time-invariant covariates. Y_{it} is one of our outcomes of interest. In the baseline analysis, we employ the post-childbirth earnings, the earnings change around the birth event, and weeks worked, contract type, and qualification defined as those observed in the year following childbirth.

The β_k's are the coefficients of interest. We exclude year 2019 for multicollinearity. The β_k's with $k < 2020$ test whether mothers in jobs that are differentially exposed to remotability display heterogeneous trends in their outcomes leading up to the pandemic. The β_k's with $k \geq 2020$ instead capture the effects of an increased remotability on the outcome once WFH contracts become available.

To better interpret the magnitudes, we will rescale these coefficients to reflect a move from the 10^{th} to the 90^{th} percentile of the $\text{Remotable}_{s(i)}$ distribution. Moving from the 10^{th} to the 90^{th} percentile in the instrument corresponds, holding broadband fixed, to shifting from highly remotable activities (e.g., web portal activities) to low-remotability sectors (e.g., aircraft manufacturing), and, holding remotability fixed at a high level, to moving from high-speed areas (e.g., Milan) to low-speed ones (e.g., Brindisi).

To quantify the effect of holding a work-from-home contract, we run the following

$$Y_{it} = \alpha_{s(i)} + \lambda_t + \sum_{k=2016}^{2024} \beta_k \cdot I\left(t = k\right) \cdot \text{WFH}_i + \sum_{k=2016}^{2024} \gamma_k \cdot I\left(t = k\right) \cdot X_i + \varepsilon_{it} \quad (2)$$

where we instrument the likelihood of holding a WFH contract (WFH_i) with the sector-location exposure to remotable work ($\text{Remotable}_{s(i)}$).

Identification Assumptions. Our identification hinges on the usual parallel trends assumption. We require that the outcomes of mothers in sector-locations differentially exposed to remotability would proceed in parallel absent the pandemic shock that boosted the introduction of WFH contracts. The confounders to our identification are time varying unobservables that correlate with $\text{Remotable}_{s(i)}$. We can classify such confounders into two macro-categories. To discuss them, without loss of generality, we rewrite the error term as $\varepsilon_{it} = \eta_{s(i)t} + \nu_{it}$, where $\eta_{s(i)t}$ includes sector-location time-varying unobservables and ν_{it} includes individual-level time-varying unobservables.

Time-varying sector-location shocks $\eta_{s(i)t}$ that could affect mothers' outcomes and are correlated with the $\text{Remotable}_{s(i)}$ would bias our estimates. These include secular employment and wage trends that differentially affect more and less remotable sector-locations. While their presence is not directly testable, the β_k's before 2020 in equation (1) check whether the parallel trend held before the pandemic and provide some empirical test of the plausibility of the absence of secular differential labor market trends across more or less remotable sector-locations. An alternative set of confounders within $\eta_{s(i)t}$ is composed of sector-location shocks that occur simultaneously to the pandemic. For instance, if the product demand for firms in a remotable sector-location changed more than in less remotable ones right after the COVID-19 shock, mothers' outcomes would change per effect of this demand change and our estimates would be confounded. There are at least two arguments that help us reduce this concern. The first is that we look at a medium-run horizon after the pandemic, which requires these shocks to be permanent to affect our estimates. This seems to be quite unlikely since no other work we are aware of has documented permanent sector-specific demand shocks after COVID-19 correlated to their remotability. The second argument is that we can test if it is plausible that sector-location-specific shocks simultaneous to COVID-19 have affected the performance of firms. This is what we do in Section 5.2, showing little evidence in support of this view. For these reasons, we believe the confounders in $\eta_{s(i)t}$ are not a major concern in our analysis.

The second macro-category of confounders consists of individual-specific time-varying shocks ν_{it}. We estimate our difference-in-differences on a repeated cross-section of mothers. Hence, we require that the individual-level unobservables are uncorrelated with $\text{Remotable}_{s(i)}$. This assumption is violated if the availability of flexible work arrangements affected the choice of becoming a mother. A change in the group of mothers over time would indeed confound our treatment effect. Suggestive evidence on a change in selection can be provided by investigating if the pre-childbirth employment history and observable time-invariant characteristics of mothers are affected by our instrument. In

the following paragraph, we explain how using a proxy for mothers' earnings potential to capture changes in selection allows us to correctly estimate the reduction in child penalty attributable to the availability of flexible work arrangements.

Changes in the Selection into Motherhood and Child Penalty Reduction. To isolate the reduction in the child penalty caused by the exposure to flexible work arrangements, we would ideally keep the selection into motherhood constant and quantify the change in post-childbirth outcomes. However, we cannot exclude that the expansion of WFH affects the selection into motherhood, changing the earnings potential of the employees who decide to have kids. Using an empirical framework, we disentangle the bias coming from a change in selection and isolate the reduction in the child penalty. Let Y^c be the earnings after childbirth and define $\tilde{Y}$ the earnings potential of a mother. The two are related through the following expression $Y^c = \rho\tilde{Y}$, where ρ is the child penalty. Changes in observed post-childbirth earnings can arise from two mechanisms: (i) variation in the magnitude of the child penalty, and (ii) shifts in the composition of mothers. Recovering the change in the child penalty thus requires netting out shifts in mothers' earnings potential from the observed change in post-childbirth earnings. As a first step, we use the change in earnings around childbirth—defined as the difference between post- and pre-birth earnings—as the outcome variable. This transformation differences out pre-childbirth earnings and accounts for compositional changes, providing a measure of the reduction in the child penalty expressed in earnings terms.

We next directly estimate the change in the child penalty, defined as the percentage reduction in a mother's earnings around childbirth. We proxy a mother's earnings potential using pre-birth earnings. Exploiting the relationship between Y^c and $\tilde{Y}$, we implement the following formula

$$\mathrm{d}\rho = \frac{\beta^{Y^c} - \rho\beta^{\tilde{Y}}}{\tilde{Y}}, \tag{3}$$

where β^Y is the difference-in-differences estimate of the effect of WFH on a generic outcome Y. To implement the formula, we thus need estimates of WFH on pre- and post-childbirth earnings, a calibrated value for the child penalty without WFH, and a measure of average pre-birth earnings.

Exclusion Restriction. To interpret our IV results as the effect of WFH, our instrument must satisfy the exclusion restriction. In this context, it requires that the effect of being in a remotable job affects a mother's outcomes only through her job arrangement and not through changes in her firm's or coworkers' performance. A potential violation of such assumption would occur if, for instance, the availability of remotable contracts boosted the productivity in remotable sectors and therefore affected the employment and

wages of the firms operating in these industries, in turn affecting the post-childbirth outcomes for reasons unrelated to the mother's work arrangement. While previous work has already documented very limited effects of such exposure on productivity (e.g. Basso et al., 2025), we are able to provide some additional suggestive evidence that these channels are of second-order importance. Specifically, in Section 5.2 we test whether the average wages, size and performance of the firm where a mother is employed are affected by the exposure to remotable work. To do so, we replace the mother's outcomes in equation (1) with her coworkers' and firm's outcomes. We will find economically and statistically not significant results.

First stage. We assess the strength and timing of the relationship between the instrument and the probability of holding a WFH contract. To this end, we estimate equation (1), where the dependent variable is the individual probability of holding a WFH contract.

Figure 1 reveals a clear and economically meaningful first stage that emerges sharply with the diffusion of WFH. Prior to 2020, coefficients are close to zero and statistically insignificant, consistent with the limited use and reporting of WFH contracts in the pre-pandemic period. Starting in 2020, the effect of remotability becomes positive and precisely estimated, indicating that workers in more remotable sectors experience a 13-14 percentage points increase in the probability of holding a WFH contract. The effect only marginally attenuates as the use of WFH stabilizes in subsequent years, supporting the relevance of the instrument.

5 Flexible Work Arrangements and Reduction of the Child Penalty

5.1 Effects on Earnings, Employment, and Promotions

We begin by estimating specification (1) on mothers' earnings in the year following childbirth. Figure 2, Panel A, plots the estimated coefficients and Table 2 summarizes the results for this and all other outcomes we will discuss in this Section by reporting the estimates of a static difference-in-differences regression (i.e., replacing year dummies in equation (1) with an indicator for the years after 2019). To facilitate the interpretation, we report the effect of moving from the 10^{th} to the 90^{th} percentile of the exposure distribution in the sample of mothers (*i.e.*, Remotable_s). We find that greater availability of flexible work arrangements has a positive and growing effect over time. The magnitude is economically meaningful: between 2021 and 2023, post-childbirth earnings are higher by 1,100-1,300 euro in sector–location cells more exposed to remote work. This pattern is consistent with a relaxation of time-related constraints that bind particularly after

childbirth. Remote work reduces commuting and coordination costs and increases effective control over work schedules, thereby raising the feasible set of labor supply choices available to new mothers. The temporal pattern of coefficients reflects the timing of the shock. Because COVID-19 struck in the first quarter of 2020, it could only affect childbirths occurring in the final months of that year.

As discussed in Section 4, part of the increase in earnings may reflect changes in the composition of mothers giving birth during the period. To address this concern, we subtract from the outcome the pre-birth earnings to account for potential shifts in mothers' earnings potential after the pandemic. When we use as the outcome the change in earnings around childbirth—defined as the difference between earnings in the year after and two years[4] before birth—we continue to find positive effects, though the estimated coefficients are smaller than those for post-birth earnings alone. On average in the post-2019 period, the estimated earnings change goes from 411 euro in 2021 to 1,005 euro in 2023, indicating that WFH opportunities substantially mitigate the child penalty. The smaller magnitude relative to post-birth earnings suggests that pre-birth earnings also increased in more exposed sector–locations, consistent with a compositional shift toward higher-earning mothers after the pandemic.

The increase in earnings is mirrored by a rise in weeks worked (Figure 2, Panel B). Mothers in more flexible sectors work about 0.7-0.8 additional weeks in the year following childbirth relative to their pre-birth levels. When we net out the change in weeks worked before childbirth, the effect remains positive and larger—around 1.7 additional week in 2023. This indicates that compositional changes in the pool of mothers do not drive the observed increase in post-birth intensive margin of labor. The fact that earnings gains are aligned with increases in weeks worked suggests that flexible work primarily operates through the labor supply margin rather than through wage-setting or bargaining channels. In other words, remote arrangements appear to expand the feasible allocation of time between market work and childcare more than altering pay conditional on hours. The increase in weeks worked is also mirrored by a decline in the use of parental leave, which decreases by about 0.5-0.7 weeks after 2020. This brings further evidence that mothers were facing a time constraint.

Turning to the extensive margin, we find little effect on the probability of employment (Panel C). Most mothers remain attached to the labor force regardless of flexible work availability. However, we uncover other extensive margin responses. We find that the likelihood of being employed full-year is 2.1 percentage points higher in 2023, which is

[4]For mothers, we use earnings measured two years before childbirth as the reference period, rather than earnings in the year immediately preceding birth. Earnings in the year before childbirth may already reflect maternity leave or pregnancy-related absences. By contrast, earnings two years prior are more likely to capture regular labor market conditions and thus provide a cleaner measure of baseline earnings.

equivalent to a 6.4% increase relative to the pre-COVID average.[5] Moreover, the probability of holding a part-time contract decreases by 1.4-2.3 percentage points for the more exposed mothers (3.1-5.1% of the pre-2020 average), suggesting that mothers substitute toward full-time work when flexible arrangements are available. This substitution away from part-time work is particularly informative. It indicates that, in the absence of flexibility, part-time arrangements may function as a second-best adjustment to rigid full-time norms. When remote work becomes available, mothers are able to maintain full-time attachment while accommodating childcare demands, consistent with a reduction in employer-imposed rigidity rather than a change in preferences over work intensity.

Finally, we examine whether the benefits extend to occupational advancement. The "glass ceiling" is a well-documented driver of the gender pay gap (Albrecht et al., 2003; Arulampalam et al., 2007; Bertrand et al., 2018), and flexible work could plausibly alter mothers' access to higher-level positions. Figure 2, Panel D, shows that promotions to managerial positions in the year after childbirth increase by 0.3-0.5 percentage points, corresponding to approximately three times the pre-2020 promotion rate among new mothers. These result suggests that flexibility does not only affect contemporaneous labor supply but may also mitigate dynamic career penalties associated with childbirth. By reducing the need for discrete reductions in hours or job transitions after birth, remote arrangements may preserve continuity in career paths and lower the probability that motherhood leads to persistent downward adjustments in occupational status.

Quantifying the Effects of Work-From-Home Contracts. We next quantify the causal effect of holding a WFH contract by implementing an instrumental variables strategy, where we instrument the probability of holding such a contract with our exposure measure. Beyond estimating the impact of a specific work arrangement, this strategy allows us to assess how much of the child penalty can be attributed to rigid job structures that limit flexibility after childbirth. We estimate a more compact version of equation (1) where we interact the instrument and the endogenous variable with a dummy for the post-COVID period. Table 2 reports the reduced form and 2SLS results in panels A and B, respectively. The first stage is strong: moving from the 10th to the 90th percentile of the Remotable_s distribution raises the probability of holding a WFH contract by 13 percentage points. The corresponding KP F-statistic is 360, indicating that the instrument is highly relevant.

[5] As is standard in the literature (see, e.g., Blau and Kahn, 2017), we define a full-year employee as an individual working at least a minimum number of weeks. Specifically, we set the threshold at 36 full-time-equivalent weeks, corresponding to three quarters of work. This choice reflects the institutional setting: following childbirth, mothers are entitled to mandatory maternity leave (up to five months) and to voluntary parental leave (up to six months per parent). In 2023, the average duration of parental leave taken by mothers was approximately one quarter. Our definition therefore accounts for the fact that mothers may be on leave for part of the year while still maintaining a strong attachment to the labor market.

The IV estimates point to sizable effects. Having a WFH contract increases mothers' earnings after childbirth by almost 7,500 euro. As discussed in Section 5.1, the composition of mothers changes after COVID-19. To account for this, we focus on earnings net of the pre-childbirth level. When using the change in earnings around childbirth as the outcome, the effect amounts to about 4,400 euro relative to an average drop of 5,000 euro and a pre-birth average of 16,400 euro. This implies that holding a WFH contract reduces the post-birth earnings drop by 87%, corresponding to 27% of the pre-birth earnings. This reduction is economically large and highlights the role of flexible work arrangements in mitigating the child penalty.

Conceptually, the IV strategy identifies the effect of flexibility for mothers whose contract status changes in response to sectoral exposure to remote work—that is, for those at the margin of adopting WFH arrangements. It does not capture the effect for mothers whose jobs are inherently non-remotable, nor does it eliminate other sources of rigidity beyond the remote margin. In this sense, the estimate reflects the impact of a partial relaxation of contractual constraints within existing employment relationships. From this perspective, the magnitude of the effect is particularly striking: if relaxing a single contractual dimension offsets most of the observed post-birth earnings decline, this suggests that rigid job structures account for a substantial portion of the child penalty.

Estimating the Reduction in the Child Penalty. We next translate the IV estimates into a measure of the change in the child penalty, using the formula in equation (3). To do so, we take mothers' pre-birth earnings as a proxy for their earnings potential. The key parameter, ρ, is calibrated in two alternative ways. In our baseline, we follow Casarico and Lattanzio (2023) and Kleven et al. (2019) to estimate a child penalty for our sample of mothers. Specifically, for each mother, we focus on the first childbirth and match her to a placebo mother defined as a woman without children who shares similar observable characteristics. We construct a balanced panel of earnings, assigning zero earnings in years when a woman is not observed in the data, and use as the outcome annual earnings normalized by earnings in the year prior to childbirth. Figure A.1 shows that in the first year after childbirth the penalty is $\rho = 0.497$. As a robustness exercise, we also compute ρ as the ratio between the pre- vs. post-childbirth change in earnings and pre-birth earnings, which yields $\rho = 0.728$. Finally, we use average pre-birth earnings for women in the post-COVID period as the denominator in equation (3).

Across both calibrations, WFH contracts lead to a sizable reduction in the child penalty. When we estimate ρ following Casarico and Lattanzio (2023), the reduction amounts to 35.6 percentage points, which corresponds to a 72% decline in the penalty. Using instead the ρ based on the average earnings drop around childbirth, the penalty falls by 31.3 percentage points, or about 43%.

5.2 Exclusion Restriction: Effects on Firms and Coworkers

For our IV estimates to capture the causal effect of a WFH contract, the exclusion restriction must hold. In particular, exposure to flexible arrangements should affect mothers' outcomes only through their contractual agreement. A potential concern is that the introduction of flexible arrangements could alter the firm productivity directly, thereby affecting earnings through channels other than the contract itself.

To examine this, we estimate equation (1) on changes in firm- and coworker-level outcomes between the year preceding and the year following childbirth, with results summarized in Figure 3. Outcomes are standardized to facilitate comparisons across specifications. To measure firm performance, we consider total employment, sales, and value added—both in levels and per-worker terms. Across all outcomes, we find no statistically significant effects: all point estimates are not significant and lie within 5% of a standard deviation, indicating economically negligible magnitudes.[6] This evidence is reassuring since improved firm performance is unlikely to explain the earnings gains we document.

Despite firm's performance seems not to be impacted, the diffusion of flexible arrangements might have changed the employees' bargaining power raising wages within the firm. To assess this, we study changes in the earnings of the coworkers of the mothers in our sample, focusing on all and male coworkers. We find no effects. This suggests that the availability of flexible contracts does not shift the average earnings profile within firms, and that our results are not driven by changes in firm-level rent sharing. The fact that the earnings of mothers' taking up WFH contracts increase relative to their coworkers suggests that we are capturing the effects of the contractual arrangement rather than more general increases in workers' bargaining power.

Overall, the evidence in Figure 3 supports the interpretation that spillovers through firm performance or coworker wages do not confound our IV estimates.

5.3 Heterogeneous Effects Across Age, Earnings Distribution, and Commuter Status

If flexible work arrangements primarily relax time constraints around childbirth, their effects should be larger for the groups for whom such constraints are more binding. To assess this prediction, Table 3 studies heterogeneity in the change in earnings between the year after childbirth and two years before childbirth. In Appendix Table A.2, we complement this evidence with a more granular heterogeneity analysis that breaks groups into finer categories and reports the corresponding reduced-form and first-stage estimates.

[6] This aligns with Basso et al. (2025) who also finds limited impact of flexible arrangements on firm performance, except for a narrow set of firms. Our estimates, however, are not directly comparable, since we capture effects at the firms where mothers are employed rather than the average firm in the economy.

Panel A studies heterogeneity by age. The conceptual framework suggests stronger effects for younger mothers, whose labor market attachment is typically less established and for whom rigid schedules may be particularly costly around childbirth. The estimates are consistent with this prediction. In the reduced form, the absolute earnings gains are of similar magnitude across age groups: mothers below 30 gain about 640 euro, compared with about 770 euro for mothers aged 36 and above. However, the first stage is markedly smaller for the youngest group. This means that younger mothers obtain similar earnings gains despite a smaller increase in WFH take-up, implying a substantially larger 2SLS effect for this group. Relative to pre-birth earnings, the gains are also by far largest for younger mothers: the difference between mothers below 30 and those aged 36 and above amounts to 42 percentage points in the IV and is statistically significant in both the IV and the reduced-form. As our framework suggests, if younger mothers obtain similar gains in level despite starting from lower earnings, these gains must come from larger adjustments in hours worked. This suggests that flexibility relaxes younger mothers' constraints on some discrete labor supply margins, such as part-time work and employment attachment.

Panel B turns to heterogeneity by the mother's pre-childbirth earnings. In absolute terms, the reduced-form effect is much larger below the median than above it, and the same is true in the 2SLS estimates. Relative to pre-birth earnings, the contrast is even sharper: the percentage gains are far larger for mothers below the median. The first stage is positive for both groups and, if anything, stronger for mothers above the median. Hence the larger effects for lower earners are not driven by a stronger take-up response to exposure. Rather, they suggest that a given increase in flexibility translates into larger labor-market gains for mothers with lower baseline earnings. As in the conceptual framework, this is consistent with lower-earning mothers responding to childbirth along more discrete margins of labor supply, such as part-time work or temporary withdrawal from employment, margins that flexible arrangements help relax.

Panel C considers heterogeneity by the father's pre-childbirth earnings. If household-level time and budget constraints are more severe in lower-income households, one would expect mothers to benefit more from flexibility when the father is a lower earner. The estimates point in that direction. The reduced-form effects are very similar across the two groups, whereas the 2SLS effect is larger when the father is below the median, about 4,700 euro compared with about 3,400 euro above the median. The first stage is positive in both groups and somewhat stronger when the father is above the median. As in the previous panel, this means that the larger IV effect for lower-income households is not driven by a stronger first stage. At the same time, the difference in coefficients across the two groups is imprecisely estimated in levels, but statistically significant when expressed in percentage points of pre-childbirth earnings. We view this as suggestive evidence that flexibility is especially valuable in households with lower male earnings and that therefore the child penalty's drivers operate at the household level.

Panel D studies commuting status. If the child penalty reflects a binding time constraint, then one should expect stronger effects for mothers who had to commute before childbirth, since commuting absorbs time that flexible work can directly save. Our estimates are consistent with this mechanism. We define the commuter status based on whether a mother works in a municipality different from the one of residence. Both the reduced-form and 2SLS effects are larger for commuters than for non-commuters: the reduced-form effect is about 670 euro for commuters versus about 385 euro for non-commuters, while the 2SLS effect is about 4,900 versus 3,500 euro. The first stage is positive and very similar across the two groups, if anything slightly smaller for commuters. This again suggests that the heterogeneity mainly reflects differences in the payoff from flexibility rather than in the induced probability of holding a WFH contract. The difference in levels is statistically significant in both the reduced form and the IV specification, although the percentage-scaled difference is not precisely estimated. Overall, these results are in line with the idea that commuting makes post-childbirth time constraints more binding and therefore increases the cost of childbirths.

Taken together, the heterogeneity results point to a common conclusion. Flexible work arrangements generate the largest gains precisely for the mothers for whom the labor supply costs of rigid workplace organization appear most severe: younger mothers, lower-earning mothers, and mothers who face commuting costs. The results by father's earnings further supports the view that the relevant margin is not only individual but household-level. These findings are therefore consistent with the conceptual framework in which time and organizational constraints are a key driver of the labor cost of motherhood.

6 Household-Level Mechanisms: The Role of Fathers' Remotability

6.1 Fathers' Labor Market Outcomes

As a first step toward understanding the household-level consequences of the expansion of remotable contracts, we examine whether fathers' labor market outcomes respond to increased exposure to flexible work arrangements around childbirth. Using the administrative link between mothers and fathers, we assign each father an exposure measure based on the industry-location in which he works at the time of childbirth and implement the same difference-in-differences strategy used for mothers.

We focus on the change in fathers' earnings around childbirth, reported in Figure 4. The figure reports effects for all fathers and for those in households where both partners are employed at the time of childbirth (the corresponding first stage estimates are shown in Figure A.3, blue line). In contrast to mothers, we find no significant effect of

exposure to remotable work on fathers' earnings. Fathers in more exposed sector-locations do not experience any differential change in earnings in the year following childbirth relative to less exposed fathers—except for a marginally significant effect in 2021 for the sample of fathers in households. This null result is consistent with the extensive literature documenting the absence of a child penalty for fathers: since fathers' earnings do not decline at childbirth, there is no direct penalty for flexible work to mitigate.

The absence of earnings responses among fathers is informative for the mechanism through which flexible work arrangements can impact the costs of motherhood for the household. It suggests that any impact of fathers' remotability on household outcomes would not operate through increases in fathers' labor supply or earnings, but rather through changes in intra-household time allocation.

6.2 Fathers' Remotability and Mothers' Child Penalty

We next turn to a joint analysis of mothers and fathers. Specifically, we investigate whether the child penalty faced by mothers is affected by the remotability of their partner's job. To do so, we augment our baseline specification for mothers' outcomes by including both mothers' and fathers' exposure to remotable work.

The results reported in Figure 5 reveal that both exposures matter. As in our baseline analysis, a mother's own exposure to remotable work increases her change in earnings (panel A) and weeks worked (panel B) around childbirth, though the magnitude is slightly smaller once father exposure is included. We find an improvement up to 700 euro (equivalent to 4% of pre-childbirth earnings). The attenuation reflects the fact that fathers' exposure independently improves mothers' outcomes and is positively correlated with mother's exposure. We find that fathers' exposure to flexible arrangements increases mothers' earnings around childbirth up to 600 euro (3.4%)—a magnitude similar to the direct effect of mothers' own remotability.[7]

These results provide strong evidence in favor of household-level constraints as a central mechanism for the labor cost of motherhood. A purely mother-side explanation—such as mother-specific career ladders, contracting frictions, or discrimination—would predict a first-order role for mothers' own flexibility and only limited effects of fathers' flexibility. If the binding constraint were located solely in mothers' jobs, relaxing fathers' constraints would not substantially alter mothers' post-childbirth earnings.

By contrast, comparable magnitudes arise naturally if the child penalty reflects a household time feasibility constraint interacting with rigid workplace arrangements for *both* spouses. If fathers' jobs are inflexible prior to the diffusion of WFH, childbirth generates a binding household time constraint that is disproportionately absorbed by mothers.

[7] Figure A.3 shows that the first-stage for fathers is similar to the one of mothers, so reduced-form estimates can almost be compared one-to-one.

When fathers' flexibility increases, they can reallocate time toward childcare—likely out of leisure, consistent with the absence of fathers' earnings responses. This relaxes mothers' time constraints and raises mothers' earnings. In this framework, flexibility on either side of the household relaxes the same underlying constraint, producing similar effects of mothers' and fathers' remotability on the child penalty.

Taken together, these results suggest that the child penalty is not solely a mother-specific labor market phenomenon but rather the outcome of joint household constraints interacting with rigid workplace arrangements. The diffusion of flexible work appears to reshape intra-household time allocation in ways that meaningfully affect mothers' earnings trajectories.

7 Flexible Arrangements and Fertility

We next examine whether the diffusion of flexible work arrangements affected fertility decisions among working women. We test whether increased availability of remote work led to higher childbirth rates, first focusing on the intensive and extensive margins of fertility among working women, and then quantifying aggregate effects at the sector–location level.

Micro-level Evidence. Our conceptual framework links fertility to the child penalty: if flexibility reduces the labor market cost of motherhood by relaxing time and coordination constraints, it lowers the marginal cost of an additional child. The effect should be strongest on the intensive margin—higher-order fertility—where the decision depends primarily on whether existing arrangements can absorb additional childcare needs. For first-time mothers, the decision also involves a fundamental identity transition shaped by partner availability and preferences for motherhood, factors unlikely to respond to changes in workplace organization alone. Accordingly, we expect the fertility response to be stronger at higher parity than on the extensive margin of entry into motherhood, unless flexible arrangements are also sufficient to offset the fixed costs of having a first child.

We first focus on the intensive margin and construct two samples. The first (*placebo*) includes women working in 2015 who, by end of 2016, had at least one child aged five or younger. The second includes women working in 2019 who, by 2020, had at least one child aged five or younger. For each woman, we build a balanced panel tracking fertility outcomes over the four years following her baseline year, pool the two panels, and estimate:

$$F_{it} = \alpha_{s(i)} + \lambda_t + \sum_{k=2015}^{2024} \beta_k \cdot I\left(t = k\right) \cdot \text{Remotable}_{s(i)} + \sum_{k=2015}^{2024} \gamma_k \cdot I\left(t = k\right) \cdot X_i + \varepsilon_{ist}, \quad (4)$$

where F_{it} is residual cumulative fertility of woman i in year t, obtained by subtracting a linear time trend in the variable $\text{Remotable}_{s(i)}$ fitted in the pre-period. $\text{Remotable}_{s(i)}$ is our remotability measure based on the industry–location of employment in the year prior to the baseline year, and X_i includes two-digit sector fixed effects. We use 2020 as the reference year, since fertility outcomes for the 2017–2020 sample are predetermined with respect to the COVID shock at that date.

We find no evidence that fertility dynamics in the placebo sample (baseline 2016) are predicted by sector-level remotability (Figure 6 Panel A). In contrast, after the expansion of remote work, fertility rises among women in more remotable sector-locations. On the *intensive margin*, births increase sharply in 2021: moving from the 10th to the 90th percentile of remotability raises the cumulative number of children by 0.013-0.014, corresponding to 1.83% of the average number of cumulative births in the placebo sample. The absolute effect stays approximately constant over the four years of the panel.

We obtain qualitatively similar results on the *extensive margin*. We define the two panels analogously, but restricting attention to women who did not have any children in the year prior to the baseline year. We again find no pre-trend in the placebo sample and a positive effect of remotability on fertility after the COVID shock (Panel B): first births rise by 0.005 in 2021 (0.2% of the placebo average) and the effect declines over the following four years. That flexible arrangements raise even first-birth rates suggests that their gains are sufficient to compensate part of the fixed costs of entering parenthood. Nevertheless, the smaller magnitude relative to the intensive-margin estimates indicates that these fixed costs are not absorbed for all households, consistent with the remaining role of anticipatory career concerns and other non-economic factors relevant for the entry into motherhood.

These findings reinforce the household-constraint interpretation: pre-COVID rigidities were distorting family choices, not only post-birth earnings trajectories.

Aggregate Fertility Effects. We next quantify the aggregate fertility response by adapting our difference-in-differences specification in (1) to the sector–location level. For each sector–location–year cell, we compute the number of mothers, construct a balanced panel including zeros for unobserved combinations, and retain the same treatment variable.

Figure 6 Panel C reports results also distinguishing by age group: women below 26, aged 27–35, and above 35. We find positive and statistically significant effects on the total number of mothers: moving from the 10^{th} to the 90^{th} percentile of the Remotable_s distribution increases the number of mothers in a location by 0.53, corresponding to 1.2% of the pre-2020 average. The effect in levels is driven by the two older groups: the number of mothers aged 27–35 increases by 0.26 and that of mothers aged 35+ by 0.24—relative to their respective pre-COVID averages of 20.90 and 20.44. These results suggest that some

women in older age groups chose to have children following the expansion of flexible work, contributing to the compositional shift in mothers and the earnings increase documented in Section 5.1.

Whether this fertility increase represents a timing adjustment or a structural rise in total fertility cannot be determined with data extending only a few years after the COVID shock. Even if the effect reflects a mere anticipation of births, however, it implies that many women had been postponing motherhood in the absence of flexible arrangements, indicating the importance of workplace organization for household well-being and female labor market dynamics.

8 Flexible Arrangements and the Gender Earnings Gap

Having established that flexibility in work arrangements is a central driver of child penalties, we conclude by quantifying how the diffusion of WFH can affect the gender earnings gap. A well-documented fact is that this gap widens over the course of a cohort's lifetime (Arellano-Bover et al., 2024; Goldin, 2014). Among several contributing factors, the child penalty is one of the most important (Goldin et al., 2024).

Using a back-of-the-envelope exercise based on our estimates, we assess the extent to which work-from-home opportunities may reduce this widening. We focus on cohorts born between 1974 and 1979 who reached the end of their fertile years before the COVID-19 shock. From a sample of individuals born in 24 birth dates per year, we select those who worked at least 156 weeks (*i.e.*, three full years) before turning 30. We then construct a balanced panel, assigning zero earnings to years in which individuals are not observed.

As expected, the earnings gender gap widens over the lifecycle of these cohorts. Figure 7 plots the ratio of female-to-male average earnings from 25 to 40 years of age. At age 25, women earn about 89.4% of men's earnings. The ratio consistently declines to reach 70.4% at the age of 40. Overall, women's earnings fall by almost 20 percentage points relative to men's over the first 15 years of the career.

We then construct three counterfactuals for the female-to-male earnings ratio: (i) one that removes the full child penalty; (ii) one that incorporates the reduction in the child penalty under the current WFH adoption (measured in 2020);[8] and (iii) one that projects the gender gap reduction if all remotable jobs adopted WFH contracts. The first counterfactual delivers the largest reduction in the gap, while the second yields the smallest. The third, lying between the two, illustrates the potential for further gains if WFH diffusion expanded.

[8]We measure the current WFH diffusion in 2020 since it is the most recent year where we have information on 4-digit sectors in the Italian Labor Force Survey.

The results highlight two main insights. First, the child penalty accounts for virtually the entire widening of the gender earnings gap over the cohort lifecycle. Eliminating it would make the female-to-male earnings ratio roughly flat or even slightly increasing in women's late thirties. This pattern reflects the fact that women have higher college completion rates and, consequently, steeper expected earnings profiles over the life cycle. Consistent with prior research, this result confirms that the child penalty is a central driver of gender inequality over the career. Moreover, because a share of women already have children by age 25, removing the child penalty would also shift the earnings ratio upward by about two percentage points.

Second, WFH arrangements attenuate the widening of the gender earnings gap. Under current adoption rates, the female-to-male ratio improves by 2.4 percentage points relative to the observed path, equivalent to 13% of the lifetime widening. This reduction would be further amplified if adoption were to expand to all remotable jobs. In this case, the female-to-male earnings ratio at age 40 would fall by only 14.1 percentage points, corresponding to about 26% of the observed widening. These calculations suggest that flexible work has the potential to reduce the lifecycle growth of gender earnings gaps, even if it is far from eliminating them entirely.

To better understand how the diffusion of work-from-home arrangements affects the evolution of the gender earnings ratio, we examine the distribution of mothers across sectors with different degrees of remotability. In particular, we assess whether mothers in our reference cohorts were concentrated in sectors more likely to benefit from flexible work arrangements, which would have favored the estimated improvement in the lifecycle earnings gap.

Figure A.4, panel A, plots the relationship between the share of mothers (measured at age 40) in each sector and the sector's share of remotable jobs. Each sector is weighted by its number of female employees. The correlation is positive but small: moving from 0 to 100% remotable jobs in a sector is associated with a 2.2–percentage point increase in the share of mothers, compared to an average share of 68%. This suggests that mothers in our cohorts were not disproportionately concentrated in highly remotable sectors prior to the COVID-19 shock. In other words, the improvement in the lifecycle earnings gap we estimate is not driven by a preexisting concentration of mothers in sectors naturally suited to remote work.

We then investigate whether post-COVID fertility patterns differ across sectors in a way that could favor a stronger reduction in the child penalty going forward. We compute each sector's post-COVID fertility rate as the share of women who experience a birth event after the pandemic out of the total number of women employed in that sector, and relate it to the sector's share of remotable jobs. The relationship is positive and economically meaningful (Figure A.4, panel B): moving from a sector with no remotable jobs to one with 100% remotable jobs increases the fertility rate by 0.23 percentage points, against

an average rate of 0.9%, which is equivalent to a 25.5% increase in fertility and it is statistically significant at 1%. These results indicate that, unlike earlier cohorts, younger cohorts of women tend to have children in sectors more exposed to the diffusion of work-from-home arrangements. This pattern is consistent with the findings in Section 7, which showed larger fertility increases for mothers with higher WFH exposure. It follows that the increase in the female-to-male earnings ratio quantified in Figure 7 may represent a lower bound for the potential improvement that younger generations of women are likely to experience as flexible work continues to spread.

9 Conclusions

This paper provides causal evidence that rigid workplace arrangements are a key driver of the child penalty. Exploiting the sharp and heterogeneous diffusion of WFH contracts after COVID-19, and combining administrative data on the universe of WFH agreements with matched employer–employee records and mother–father links, we show that holding a WFH contract offsets a large share of the earnings loss associated with childbirth. Because our estimates capture only the margin of jobs that can become remote and only those mothers induced into flexible arrangements, they should be interpreted as a lower bound on the contribution of contractual rigidity to the child penalty.

Our results also clarify the mechanism behind the penalty. Fathers' earnings do not respond to greater exposure to remote work around childbirth, yet fathers' remotability reduces mothers' earnings losses by an amount comparable to mothers' own flexibility. This pattern is difficult to reconcile with explanations based solely on mother-specific career dynamics, discrimination, or fixed gender norms. Instead, it points to a binding household-level time constraint that interacts with rigid job structures on both sides of the couple. Flexibility relaxes this joint constraint and allows mothers to remain more strongly attached to work. Consistent with this interpretation, we find that greater availability of flexible contracts lowers the labor market cost of the first and marginal child and increases fertility among working women and mothers. In a context of declining fertility in advanced economies, this highlights how workplace organization can shape not only earnings trajectories but also family decisions. Moreover, if the objective is to promote paternal involvement without lowering fathers' earnings and to reduce the labor market cost of children for mothers, expanding access to flexible work arrangements may be more effective than paternal leave policies that entail implicit or explicit earnings penalties for fathers (Andresen and Nix, 2025).

Finally, we show that the post-COVID diffusion of WFH has the potential to meaningfully slow the lifecycle widening of the gender earnings gap. While flexible work alone cannot eliminate gender inequality, our evidence indicates that contractual rigidity is a first-order component of the child penalty and of the divergence in men's and women's

careers after parenthood.

Our evidence suggests that the child penalty is not solely driven by preferences or norms, but is shaped to a significant extent by institutional and contractual features of work. Policies and firm practices that expand time flexibility—while preserving career progression—can therefore play a central role in reducing gender inequality and reshaping labor supply and fertility decisions within households.

References

Aksoy, C. G., Barrero, J. M., Bloom, N., Cranney, K., Davis, S. J., Dolls, M., and Zarate, P. (2026). Working from Home and Fertility. Technical report, Working Paper.

Albrecht, J., Björklund, A., and Vroman, S. (2003). Is There a Glass Ceiling in Sweden? *Journal of Labor Economics*, 21(1):145–177.

Andresen, M. E. and Nix, E. (2025). You Cannot Force Me Into Caregiving: Paternity Leave and the Child Penalty. *The Economic Journal*, 136(674):780–797.

Angelov, N., Johansson, P., and Lindahl, E. (2016). Motherhood and the Female Labor Market. *Journal of Labor Economics*, 34(3):545–579.

Arellano-Bover, J., Bianchi, N., Lattanzio, S., and Paradisi, M. (2024). One Cohort at a Time: Young Workers' Outcomes and the Declining Gender Pay Gap. *Working paper*.

Arntz, M., Ben Yahmed, S., and Berlingieri, F. (2022). Working from home, hours worked and wages: Heterogeneity by gender and parenthood. *Labour Economics*, 76:102169.

Arulampalam, W., Booth, A. L., and Bryan, M. L. (2007). Is There a Glass Ceiling over Europe? Exploring the Gender Pay Gap across the Wage Distribution. *Industrial and Labor Relations Review*, 60(2):163–186.

Atkin, D., Schoar, A., and Shinde, S. (2023). Working from Home, Worker Sorting and Development. *NBER Working Paper*.

Barrero, J. M., Bloom, N., and Davis, S. J. (2023). The Evolution of Work from Home. *Journal of Economic Perspectives*, 37(4):23–50.

Basso, G., Boeri, T., Caiumi, A., and Paccagnella, M. (2022). Unsafe Jobs, Labour Market Risk and Social Protection. *Economic Policy*, 37(110):229–267.

Basso, G., Dottori, D., and Formai, S. (2025). Working from home and labour productivity: Firm-level evidence. Working Paper 1508, Banca d'Italia.

Bertrand, M., Black, S. E., Jensen, S., and Lleras-Muney, A. (2018). Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labour Market Outcomes in Norway. *The Review of Economic Studies*, 86(1):191–239.

Blau, F. D. and Kahn, L. M. (2017). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature*, 55(3):789–865.

Bloom, N., Han, R., and Liang, J. (2024). Hybrid working from home improves retention without damaging performance. *Nature*, 630:920–925.

Bloom, N., Liang, J., Roberts, J., and Ying, Z. J. (2015). Does Working from Home Work? Evidence from a Chinese Experiment. *The Quarterly Journal of Economics*, 130(1):165–218.

Boeri, F., Crescenzi, R., and Rigo, D. (2025). Work from home and Firm Productivity: The Role of ICT and Size. Working paper, CESifo.

Burdett, A., Etheridge, B., Tang, L., and Wang, Y. (2024). Worker productivity during Covid-19 and adaptation to working from home. *European Economic Review*, 167:104788.

Casarico, A. and Lattanzio, S. (2023). Behind the child penalty: understanding what contributes to the labour market costs of motherhood. *Journal of Population Economics*, 36(3):1489–1511.

Choudhury, P., Khanna, T., Makridis, C. A., and Schirmann, K. (2024). Is Hybrid Work the Best of Both Worlds? Evidence from a Field Experiment. *The Review of Economics and Statistics*, pages 1–24.

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

Crescenzi, R., Dottori, D., and Rigo, D. (2025). Work From Home, Labour Market Participation and Employment. Questioni di Economia e Finanza (Occasional Papers) 958, Banca d'Italia.

Cullen, Z. and Perez-Truglia, R. (2023). The Old Boys' Club: Schmoozing and the Gender Gap. *American Economic Review*, 113(7):1703–40.

Dingel, J. I. and Neiman, B. (2020). How Many Jobs Can Be Done at Home? *Journal of Public Economics*, 189.

Dunatchik, A., Gerson, K., Glass, J., Jacobs, J. A., and Stritzel, H. (2021). Gender, parenting, and the rise of remote work during the pandemic: Implications for domestic inequality in the United States. *Gender & Society*, 35(2):194–205.

Emanuel, N. and Harrington, E. (2024). Working Remotely? Selection, Treatment, and the Market for Remote Work. *American Economic Journal: Applied Economics*, 16(4):528–59.

Farooqi, H. (2023). Telework availability, women's labor market outcomes and fertility choices. Working Paper.

Fenizia, A. and Kirchmaier, T. (2025). Not Incentivized Yet Efficient: Working From Home in the Public Sector. Technical report, mimeo.

Gibbs, M., Mengel, F., and Siemroth, C. (2023). Work from Home and Productivity: Evidence from Personnel and Analytics Data on Information Technology Professionals. *Journal of Political Economy Microeconomics*, 1(1):7–41.

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

Goldin, C. (2022). Understanding the Economic Impact of COVID-19 on Women. *Brookings Papers on Economic Activity*, 2022(1):65–140.

Goldin, C., Kerr, S. P., and Olivetti, C. (2024). The parental pay gap over the life cycle: Children, jobs, and labor supply. *Journal of Economic Dynamics and Control*, 169:104963.

Goux, D. and Maurin, E. (2025). Sick of Working from Home? *The Economic Journal*, page ueaf025.

Harrington, E. and Kahn, L. (2025). Has the Rise of Work-from-Home Reduced the Motherhood Penalty in the Labor Market? Working Paper.

Heggeness, M. and Suri, P. (2021). Telework, Childcare, and Mothers' Labor Supply. Opportunity and Inclusive Growth Institute Working Papers 52, Federal Reserve Bank of Minneapolis.

Jack, R., Tannenbaum, D., and Timpe, B. (2025). The Parenthood Gap: Firms and Earnings Inequality after Kids. Upjohn Institute Working Paper.

Kleven, H., Landais, C., and Leite-Mariante, G. (2024a). The Child Penalty Atlas. *The Review of Economic Studies*, 92(5):3174–3207.

Kleven, H., Landais, C., Posch, J., Steinhauer, A., and Zweimüller, J. (2024b). Do Family Policies Reduce Gender Inequality? Evidence From 60 Years of Policy Experimentation. *American Economic Journal: Economic Policy*, 16(2):110–149.

Kleven, H., Landais, C., and Søgaard, J. E. (2019). Children and Gender Inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.

Lamorgese, A., Linarello, A., Schivardi, F., and Patnaik, M. (2024). Management Practices and Resilience to Shocks: Evidence from COVID-19. *Management Science*, (70(12)).

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2020). Gender Differences in Job Search: Trading off Commute against Wage. *The Quarterly Journal of Economics*, 136(1):381–426.

Lee, K. (2023). Working from Home as an Economic and Social Change: A Review. *Labour Economics*, 85:102462.

Mas, A. and Pallais, A. (2017). Valuing Alternative Work Arrangements. *American Economic Review*, 107(12):3722–3759.

Olivetti, C., Pan, J., and Petrongolo, B. (2024). The Evolution of Gender in the Labor Market. In Dustmann, C. and Lemieux, T., editors, *Handbook of Labor Economics*, volume 5, pages 619–677. Elsevier.

Olivetti, C. and Petrongolo, B. (2016). The Evolution of Gender Gaps in Industrialized Countries. *Annual Review of Economics*, 8:405–434.

Olivetti, C. and Petrongolo, B. (2017). The Economic Consequences of Family Policies: Lessons from a Century of Legislation in High-Income Countries. *Journal of Economic Perspectives*, 31(1):205–230.

Pabilonia, S. W. and Vernon, V. (2022). Telework, Wages, and Time Use in the United States. *Review of Economics of the Household*, 20(3):687–734.

Song, M. (2025). The Future of Work: Remote Opportunities and Female Labor Force Participation. Available at SSRN 5127427.

Zarate, P. (2025). Remote Work and Child Penalties. Available at SSRN 5360652.

Table 1: Characteristics of mothers and fathers

	(1)	(2)	(3)	(4)
	All		Households	
	Mothers	Fathers	Mothers	Fathers
Age	32.536	34.309	31.918	34.525
	(4.913)	(6.156)	(4.778)	(5.609)
Full-time	0.589	0.831	0.620	0.865
	(0.492)	(0.374)	(0.485)	(0.342)
Permanent	0.863	0.811	0.862	0.856
	(0.344)	(0.392)	(0.345)	(0.351)
Manager	0.020	0.033	0.020	0.043
	(0.141)	(0.179)	(0.141)	(0.202)
Weeks worked	34.714	36.643	36.129	40.140
	(19.070)	(20.350)	(18.670)	(18.610)
Earnings	16,424.570	19,980.773	17,126.258	22,780.139
	(13,074.196)	(18,131.904)	(12,751.108)	(18,555.893)
Prob. WFH	0.064	0.058	0.076	0.073
	(0.222)	(0.205)	(0.240)	(0.229)
South	0.252	0.289	0.224	0.223
	(0.434)	(0.453)	(0.417)	(0.416)
Industry	0.193	0.420	0.199	0.423
	(0.395)	(0.494)	(0.399)	(0.494)
Firm size 1–14	0.390	0.315	0.362	0.287
	(0.488)	(0.465)	(0.481)	(0.453)
Firm size 15–99	0.240	0.295	0.246	0.291
	(0.427)	(0.456)	(0.430)	(0.454)
Firm size 100+	0.370	0.390	0.392	0.421
	(0.483)	(0.488)	(0.488)	(0.494)
Coworker earnings	18,621.277	21,737.113	19,207.809	23,838.791
	(14,103.479)	(14,602.819)	(14,250.233)	(15,150.148)
Male coworker earnings	22,549.232	22,954.246	23,207.303	25,316.705
	(17,510.826)	(16,205.376)	(17,573.387)	(16,938.035)
Value added per worker	46.120	57.273	47.698	62.066
	(51.472)	(54.810)	(52.383)	(57.034)
Sales per worker	204.319	235.947	211.447	261.115
	(336.117)	(332.458)	(342.170)	(349.857)
Number of workers	796,174	938,667	428,397	428,397

Notes. The table reports the mean and standard deviation (in parentheses) of baseline characteristics for mothers and fathers in the baseline analysis. Columns (1) and (2) refer to the full sample of parents, while columns (3) and (4) restrict the sample to parents living in households in which both the mother and the father are employed at the time of childbirth. All variables are measured in the year of childbirth, with the exception of earnings and weeks worked, which are measured two years before childbirth.

Table 2: Reduced form and 2SLS static difference-in-differences

	(1) Earnings	(2) Net earnings	(3) Weeks	(4) Net weeks	(5) PL weeks	(6) Employed	(7) Employed full-year	(8) Part-time	(9) Promoted post
	Panel A: Reduced form								
Remotable × Post	978.18***	570.156***	0.528***	0.805***	-0.712***	-0.001	0.015***	-0.022***	0.003***
	(149.83)	(107.544)	(0.136)	(0.159)	(0.096)	(0.002)	(0.004)	(0.004)	(0.001)
Dep. var. mean	11395	-5006	25.6	-9.16	11.6	.86	.329	.457	.00192
Observations	770,625	770,625	770,625	770,625	770,625	770,625	770,625	661,561	661,561
	Panel B: 2SLS								
WFH × Post	7498.98***	4370.96***	4.05***	6.17***	-5.46***	-0.008	0.118***	-0.162***	0.023***
	(1106.81)	(839.51)	(1.02)	(1.20)	(0.71)	(0.016)	(0.029)	(0.027)	(0.004)
First stage	0.130***	0.130***	0.130***	0.130***	0.130***	0.130***	0.130***	0.130***	0.130***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
KP F-stat	360.364	360.364	360.364	360.364	360.364	360.364	360.364	363.513	363.513
Dep. var. mean	11395	-5006	25.6	-9.16	11.6	.86	.329	.457	.00192
Observations	770,625	770,625	770,625	770,625	770,625	770,625	770,625	661,561	661,561

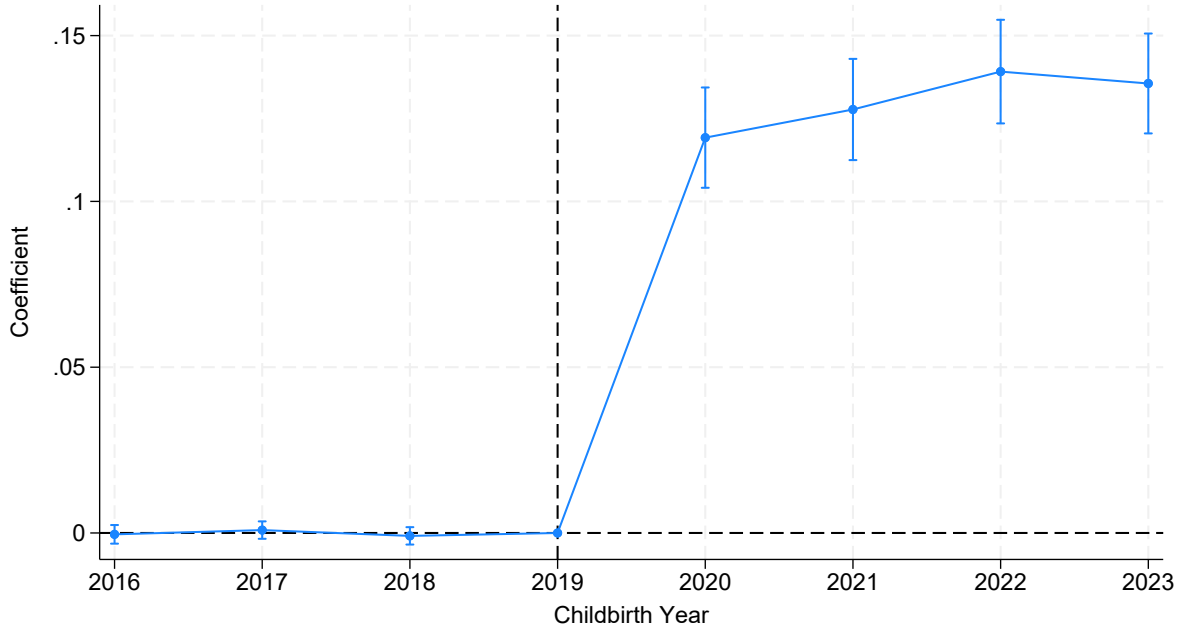
Notes. This table reports the coefficients from the estimation of equation (1) in the reduced form (Panel A) and of equation (2) in 2SLS (Panel B) in a static specification, where event-time dummies are replaced by a single post indicator equal to one for years after 2019. The reported coefficients correspond to the interaction between remotability (Panel A) or actual remote work (Panel B) and the post indicator. The outcomes, measured with a one-year lead unless otherwise specified, are: annual earnings (column 1); annual earnings net of pre-childbirth earnings (column 2); total full-time equivalent weeks worked (column 3); total full-time equivalent weeks net of pre-childbirth weeks (column 4); parental leave duration in weeks (column 5); an indicator for being employed (column 6); an indicator for being employed for at least 26 full-time equivalent weeks (column 7); an indicator for having a part-time contract, conditional on employment (column 8); an indicator for holding a managerial position after childbirth (column 9) and before childbirth (column 10); and an indicator for being promoted to a managerial position (from either blue- or white-collar jobs) (column 11). Standard errors, clustered at the 4-digit sector and local labor market level, are reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Heterogeneity: Reduced form and 2SLS static difference-in-differences

	First-stage	Reduced-form	2SLS	Pre-2020 earnings at $t-2$
Panel A: Age				
Under-30	0.06***	640.70***	7087.67***	10660.09
	(0.01)	(134.26)	(1256.85)	
30-35	0.14***	437.99***	3892.26***	16945.81
	(0.01)	(115.48)	(804.30)	
36+	0.14***	769.38***	5157.59***	21080.96
	(0.01)	(194.81)	(1130.46)	
KP F-stat	100.040			
Difference: Under-30 vs 36+		-128.677	1930.072*	
		[-0.624]	[1.748]	
Difference (%): Under-30 vs 36+		0.024*	0.420***	
		[1.740]	[4.315]	
Observations	770,625	770,625	770,625	
Panel B: Mother's earnings				
Below median	0.04***	1082.14***	13481.04***	6669.71
	(0.01)	(165.91)	(1994.47)	
Above median	0.17***	252.15**	2965.54***	25923.79
	(0.01)	(110.44)	(757.93)	
KP F-stat	109.045			
Difference: Below vs above median		829.993***	10515.500***	
		[4.944]	[6.819]	
Difference (%): Below vs above median		0.153***	1.907***	
		[6.387]	[6.841]	
Observations	770,625	770,625	770,625	
Panel C: Father's earnings				
Below median	0.09***	485.61***	4650.79***	13282.72
	(0.01)	(160.41)	(1402.02)	
Above median	0.16***	492.61***	3442.21***	21320.31
	(0.01)	(141.60)	(954.70)	
KP F-stat	152.816			
Difference: Below vs above median		-6.994	1208.579	
		[-0.049]	[1.515]	
Difference (%): Below vs above median		0.013	0.189***	
		[1.340]	[2.608]	
Observations	405,843	405,843	405,843	
Panel D: Commuting				
Commuter	0.13***	672.14***	4930.22***	17111.86
	(0.01)	(99.70)	(793.49)	
Non-commuter	0.14***	384.76**	3527.73***	15277.19
	(0.01)	(163.52)	(971.31)	
KP F-stat	179.066			
Difference: Commuter vs non-commuter		287.377*	1402.497**	
		[1.955]	[2.343]	
Difference (%): Commuter vs non-commuter		0.014	0.057	
		[1.478]	[1.455]	
Observations	770,625	770,625	770,625	

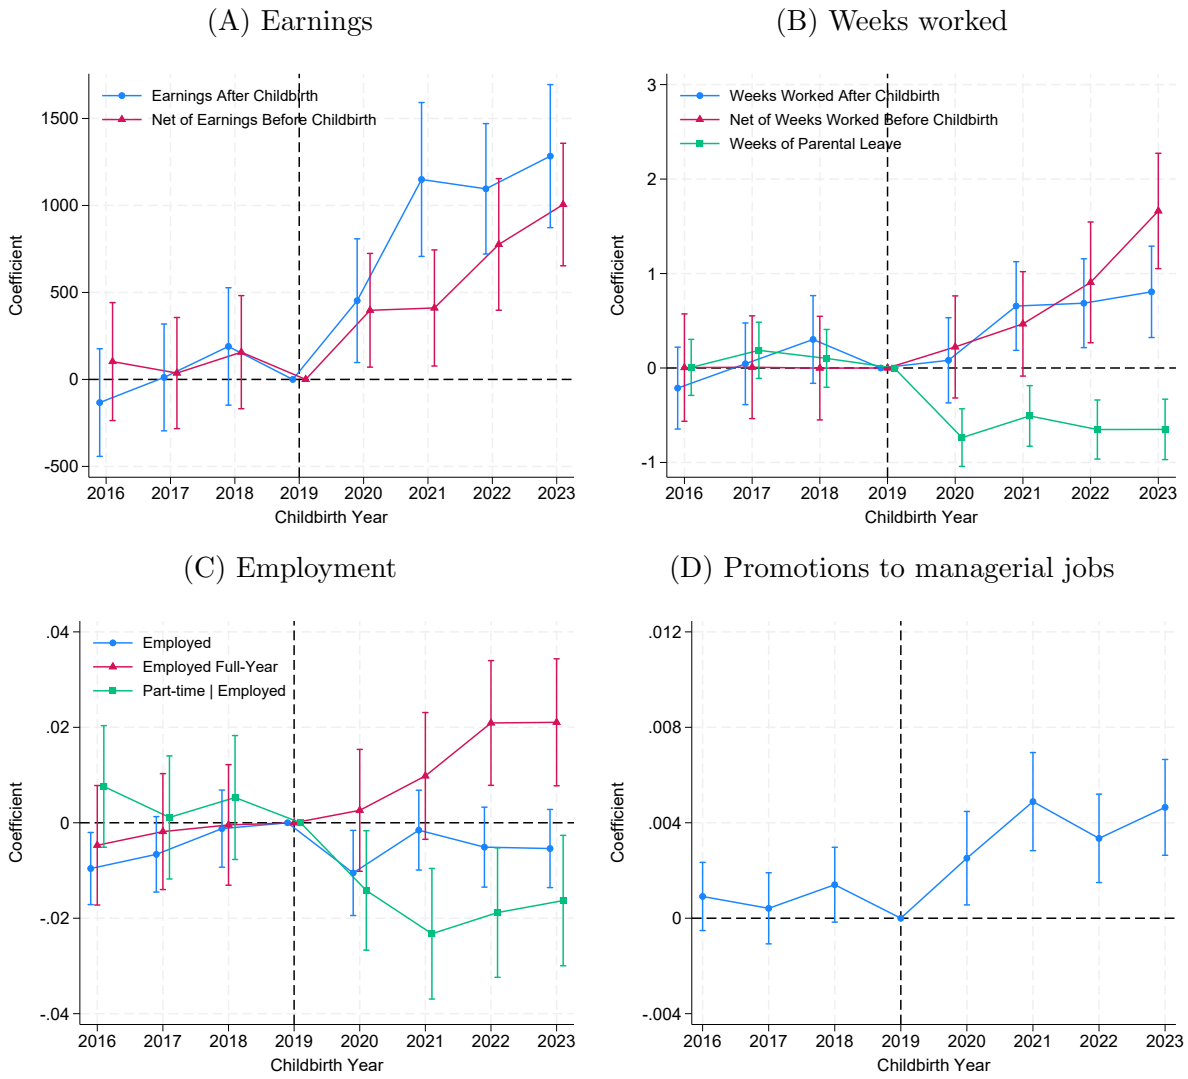
Notes. Each column reports estimates from pre-post heterogeneous DID specifications for the earnings change around childbirth. The first-stage column reports the effect of remotability on WFH, the reduced-form column reports the effect of remotability on the outcome, and the 2SLS column reports the implied causal effect of WFH. The last column reports the subgroup-specific pre-2020 mean of pre-childbirth earnings, measured two years before the childbirth year. Standard errors for subgroup coefficients are reported in parentheses, while t-statistics for difference tests are reported in square brackets. Difference (%) rescales coefficient differences by the subgroup-specific pre-2020 mean of pre-childbirth earnings. KP F-stat denotes the Kleibergen-Paap weak-instrument statistic.

Figure 1: First-stage dynamic estimates for mothers' WFH contracts



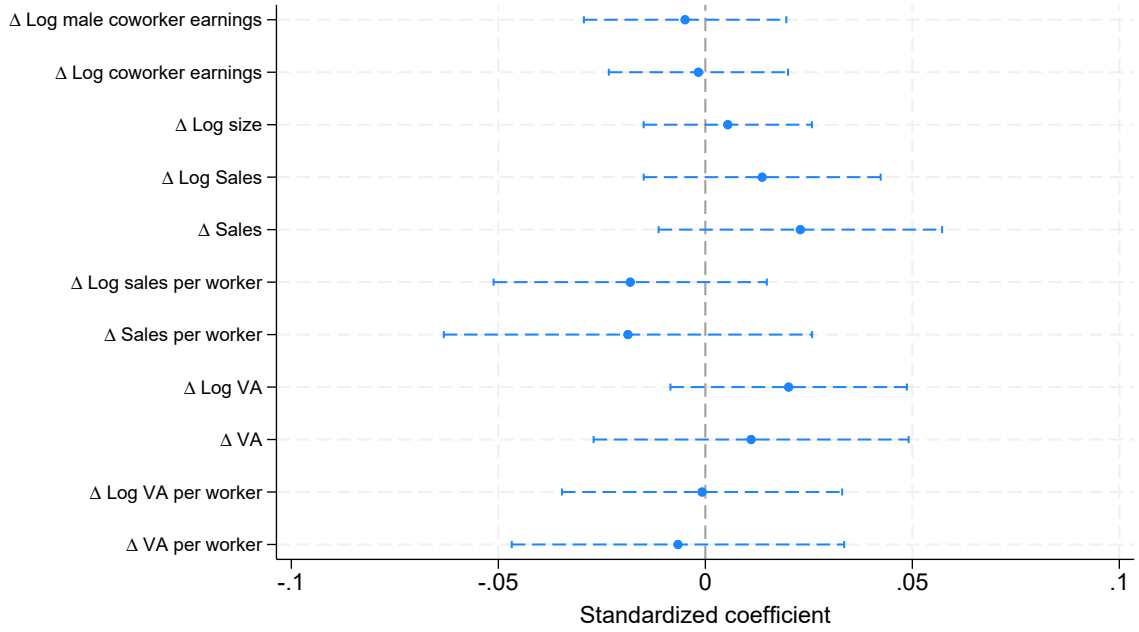
Notes. This figure plots first-stage coefficients and their 95% confidence intervals, based on standard errors clustered at the 4-digit sector and local labor market level. Coefficients come from estimating equation (1), where the outcome is the probability of holding a WFH contract.

Figure 2: Labor market outcomes of mothers, dynamic reduced form estimates



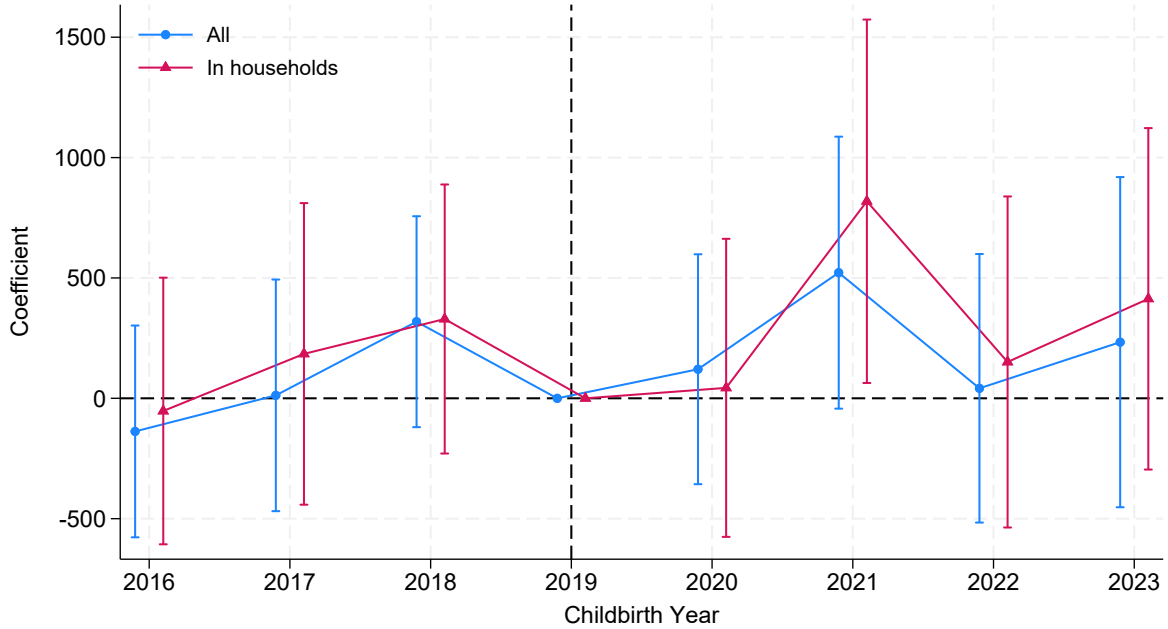
Notes. This figure plots reduced form coefficients and their 95% confidence intervals (based on standard errors clustered at the 4-digit sector and local labor market level) resulting from the estimation of equation (1) for different outcomes of interest. The outcome variables are: earnings in the year after childbirth and those netted of the pre-childbirth earnings (panel A); total full-time equivalent weeks worked in the year after childbirth, those netted of the pre-childbirth weeks and weeks of parental leave (panel B); a dummy for being employed, for being employed for at least 26 full-time equivalent weeks, and for working part-time conditional on being employed (panel C); a dummy for holding a managerial job in the year after or before childbirth and a dummy for promotion to a managerial position the year after childbirth, conditional on being employed (panel D).

Figure 3: Exclusion restriction: remotability and firm's and coworkers' outcomes



Notes. This figure plots reduced form standardized coefficients and their 95% confidence intervals (based on standard errors clustered at the 4-digit sector and local labor market level) resulting from the estimation of equation (1) for different covariates. The coefficient reported is the linear combination of the post-2019 years.

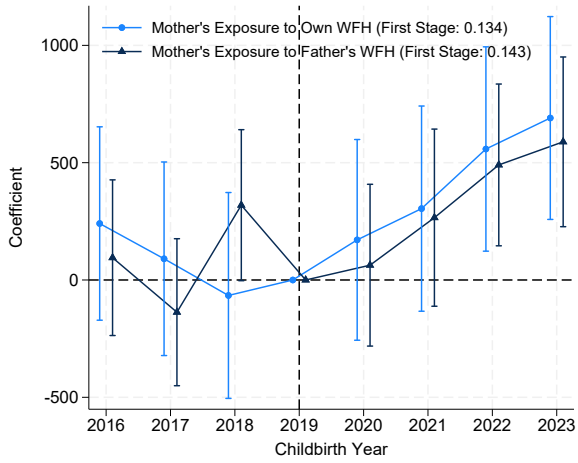
Figure 4: Earnings of fathers, dynamic reduced form estimates



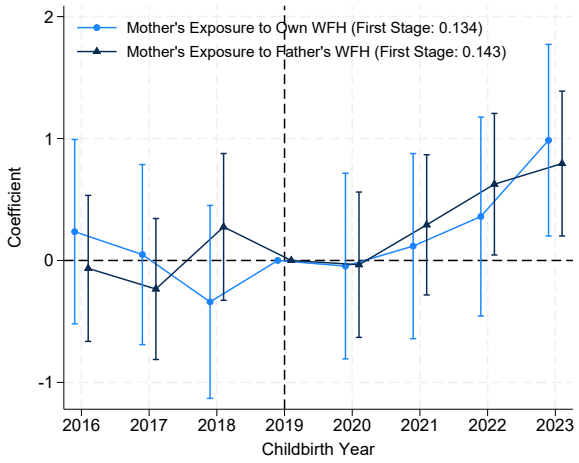
Notes. This figure reports reduced-form coefficients and their 95% confidence intervals (based on standard errors clustered at the 4-digit sector and local labor market levels) from the estimation of equation (1) for fathers. The outcome is fathers' earnings in the year following childbirth. Estimates are shown for the full sample of fathers and for the subsample of fathers living in households in which both parents are employed at the time of childbirth.

Figure 5: Earnings and weeks of mothers, exposure to both own and father's remotability, dynamic reduced-form estimates

(A) Earnings change



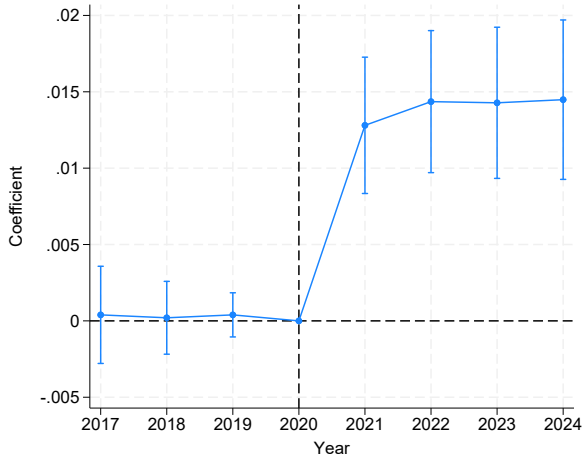
(B) Weeks change



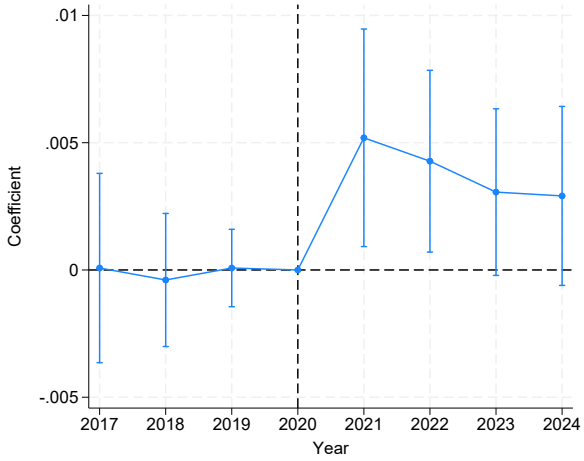
Notes. This figure plots reduced-form coefficients and their 95% confidence intervals (based on standard errors clustered at the 4-digit sector and local labor market level) from the estimation of equation (1), augmented with the remotability of the father's job. The outcomes are the change in earnings (panel A) and the change in weeks worked (panel B) between the year after and two years before childbirth, both rescaled by their respective pre-childbirth means.

Figure 6: Fertility outcomes and number of mothers, reduced-form estimates

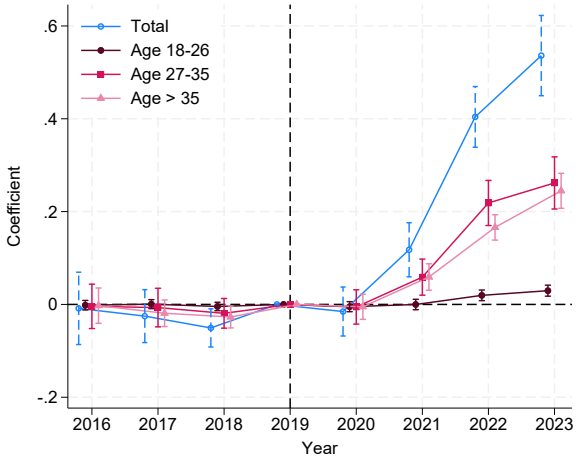
(A) Incumbent mothers, prob. second child



(B) Women, prob. first child

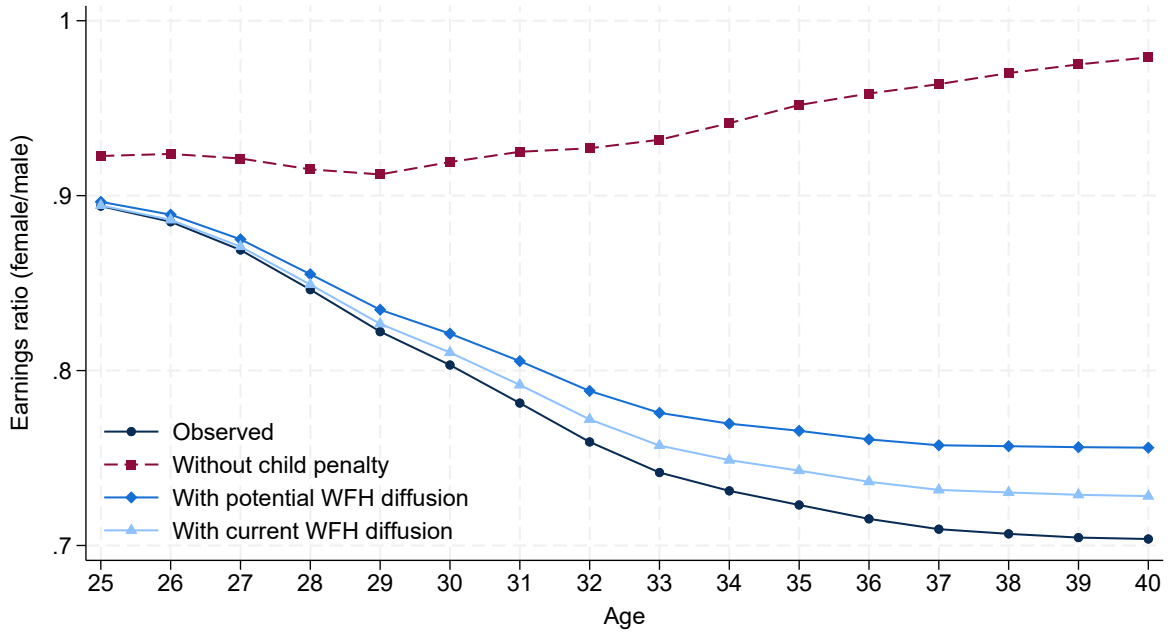


(C) Number of mothers



Notes. Panels A and B report reduced-form coefficients and their 95% confidence intervals, based on standard errors clustered at the 4-digit sector and local labor market levels, from the estimation of equation (4). Panel A is estimated on a sample of women working in 2015 who, by the end of 2016, had at least one child aged five or younger, and of women working in 2019 who, by 2020, had at least one child aged five or younger. Panel B is estimated on the corresponding sample of first-time mothers. For both groups, a balanced panel is constructed to track fertility outcomes over the five years following the baseline. The outcome in the top panels is the cumulative number of children. Panel C reports reduced-form coefficients and their 95% confidence intervals from a sector–location-level adaptation of equation (1). For each sector–location–year cell, we compute the number of mothers, construct a balanced panel including zeros for unobserved combinations, and retain the same treatment variable. The estimates are shown separately by age group: women below age 26, aged 27–35, and above age 35. All estimates are adjusted for pre-trends. Specifically, we estimate equations (1) and (4) replacing year fixed effects with a linear time trend over the pre-2020 period, and we net out the outcome by the corresponding predicted trend before estimating the main specification. Standard errors are clustered at the 4-digit sector and local labor market levels.

Figure 7: Gender gap in labor earnings over the life cycle under different child penalty scenarios



Notes. The figure shows the evolution of the gender earnings ratio over the life cycle. "Observed" is the actual ratio of female-to-male earnings between the ages of 25 and 40. "With current WFH diffusion" reports an estimate of the ratio using the average diffusion of WFH in the data. "With potential WFH diffusion" reports an estimate of the earnings ratio assuming that all workers that could potentially work remotely would do so. "Without child penalty" assumes that the child penalty in annual labor earnings is null.

Online Appendix to Remote Work and Gender Inequality after Childbirth

Gaetano Basso, Maria De Paola, Salvatore Lattanzio, Matteo Paradisi

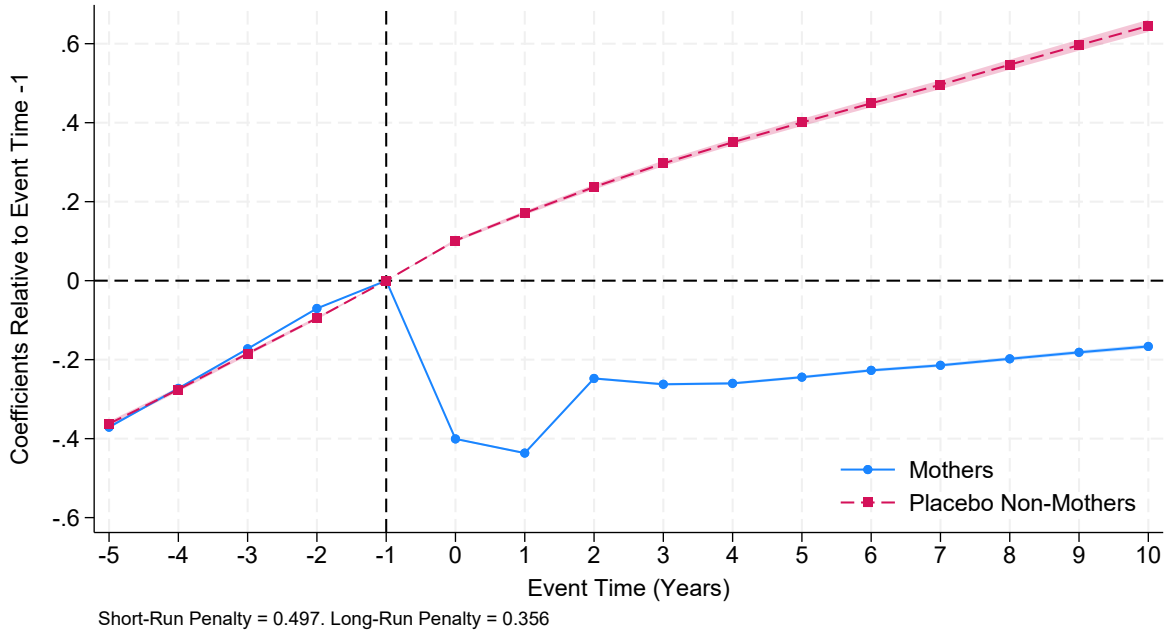
Additional Tables and Figures

Table A.1: Reduced-form static difference-in-differences: robustness

	(1) Earnings	(2) Net earnings	(3) Weeks	(4) Net weeks	(5) PL weeks	(6) Employed	(7) Employed full-year	(8) Part-time	(9) Promoted post
					Panel A: Baseline				
Remotable × Post	978.18***	570.16***	0.53***	0.81***	-0.71***	-0.001	0.015***	-0.022***	0.003***
	(149.83)	(107.54)	(0.14)	(0.16)	(0.10)	(0.002)	(0.004)	(0.004)	(0.001)
Dep. var. mean	11395	-5006	25.6	-9.16	11.6	.86	.329	.457	.00192
Observations	770,625	770,625	770,625	770,625	770,625	770,625	770,625	661,561	661,561
					Panel B: Mothers in household				
Remotable × Post	1122.75***	499.82***	0.60***	0.57***	-0.78***	-0.000	0.019***	-0.026***	0.004***
	(184.69)	(131.62)	(0.17)	(0.20)	(0.12)	(0.003)	(0.005)	(0.005)	(0.001)
Dep. var. mean	11752	-5561	26.5	-10	12.1	.867	.349	.429	.00197
Observations	405,843	405,843	405,843	405,843	405,843	405,843	405,843	350,483	350,483
					Panel C: Multi-way clustered standard errors				
Remotable × Post	978.18***	570.16***	0.53**	0.81	-0.71***	-0.001	0.015***	-0.022***	0.003
	(217.29)	(151.92)	(0.18)	(0.00)	(0.09)	(0.002)	(0.004)	(0.003)	(0.000)
Dep. var. mean	11395	-5006	25.6	-9.16	11.6	.86	.329	.457	.00192
Observations	770,625	770,625	770,625	770,625	770,625	770,625	770,625	661,561	661,561

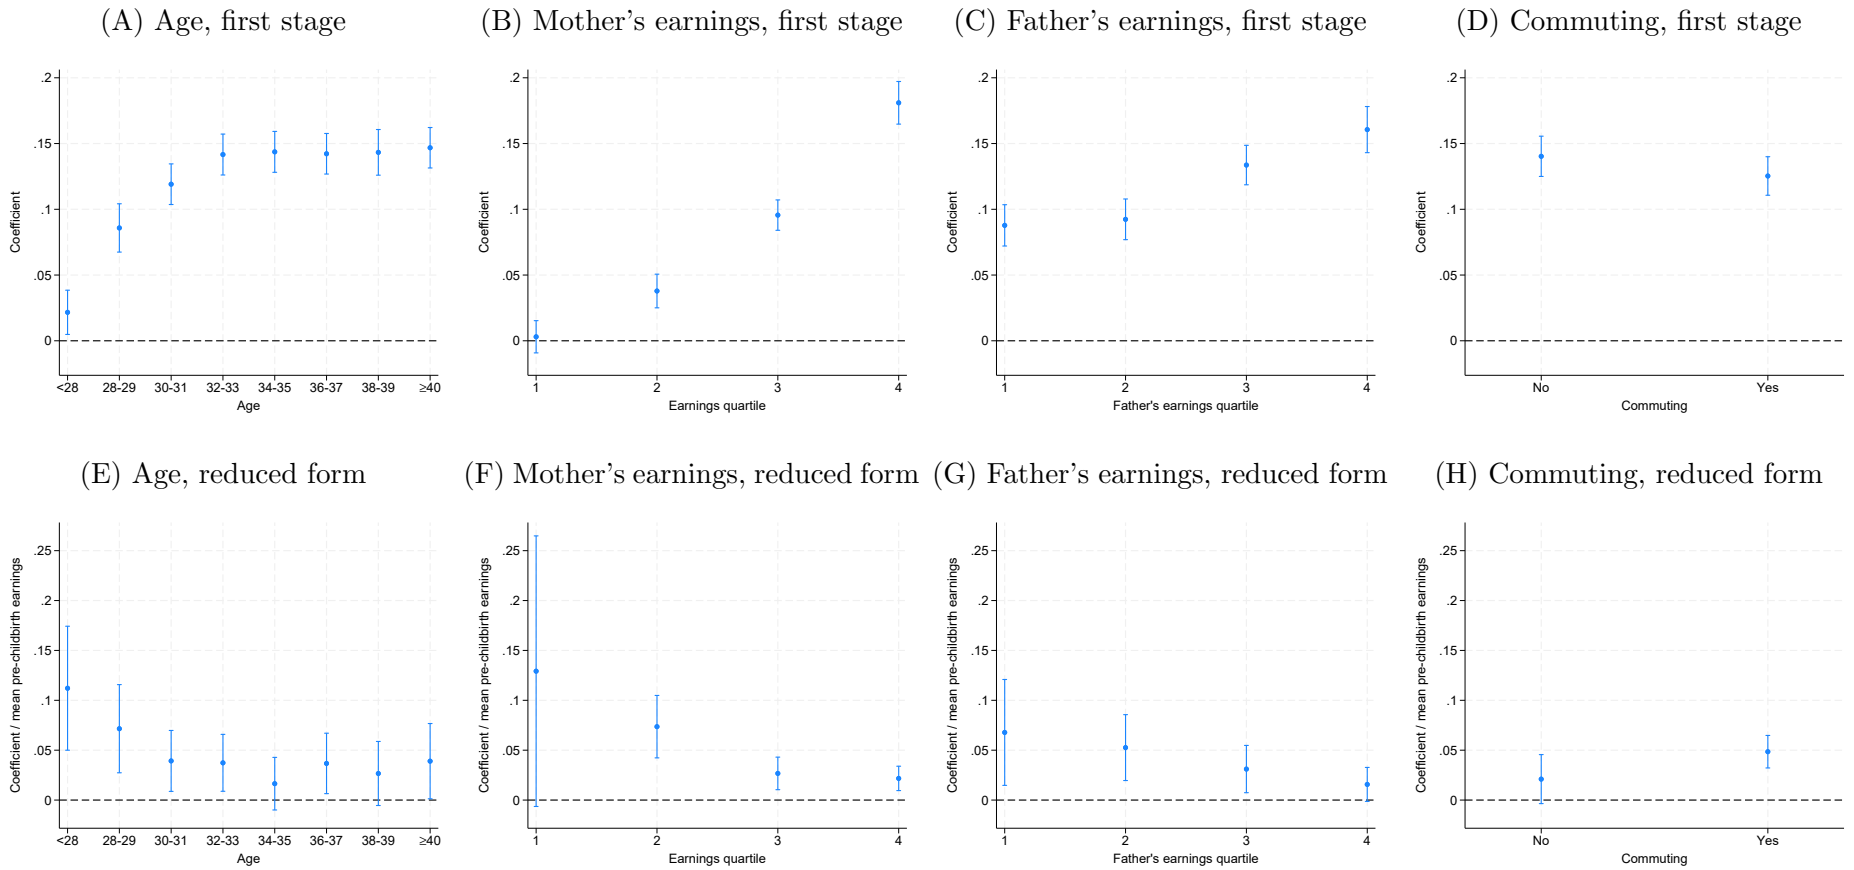
Notes. This table reports reduced-form coefficients from the estimation of equation (1) in a static specification, where event-time dummies are replaced by a single post indicator equal to one for years after 2019. The reported coefficient is that on the interaction between remotability and the post indicator. Panel A reports estimates for the baseline sample. Panel B reports estimates for the subsample of mothers living in households in which both parents are employed at the time of childbirth. Panel C reports baseline estimates with standard errors clustered multi-way at the sector, location, and year levels. The outcomes, measured with a one-year lead unless otherwise specified, are: annual earnings (column 1); annual earnings net of pre-childbirth earnings (column 2); total full-time equivalent weeks worked (column 3); total full-time equivalent weeks net of pre-childbirth weeks (column 4); parental leave duration in weeks (column 5); an indicator for being employed (column 6); an indicator for being employed for at least 26 full-time equivalent weeks (column 7); an indicator for having a part-time contract, conditional on employment (column 8); an indicator for holding a managerial position after childbirth (column 9) and before childbirth (column 10); and an indicator for being promoted to a managerial position, from either blue- or white-collar jobs (column 11). Standard errors are reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure A.1: The child penalty



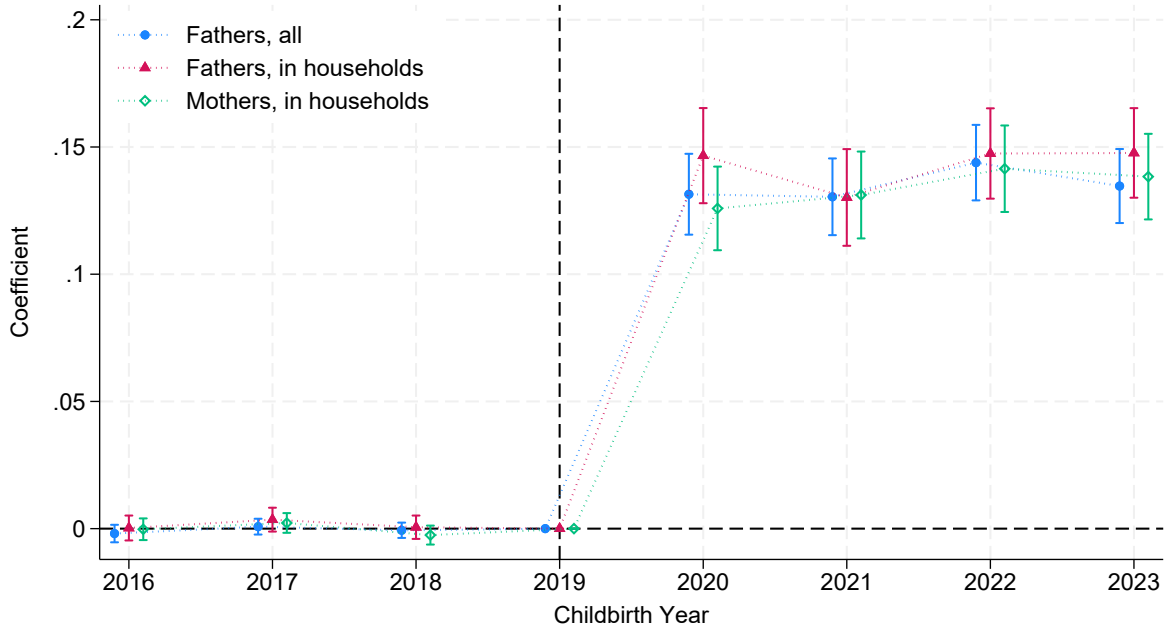
Notes. The graph displays event-time coefficients estimated from a regression of annual earnings on year-to-childbirth dummies, calendar year dummies and individual fixed effects ($Y^g_{ist} = \theta_i + \sum_{j \neq -1} \alpha^g_j \cdot \mathbf{1}[j = t] + \sum_y \gamma^g_y \cdot \mathbf{1}[y = s] + \nu^g_{ist}$), as a percentage of the counterfactual outcome in the absence of children (i.e., $P^g_t \equiv \hat{\alpha}^g_t / E[Y^g_{ist} \mid t]$) for mothers and placebo non-mothers separately. Placebo non-mothers are identified as in Casarico and Lattanzio (2023). The figure also reports a "child penalty," defined as the percentage by which women fall behind women without children after having children: $P_t \equiv (\hat{\alpha}^m_t - \hat{\alpha}^w_t)/E[Y^w_{ist} \mid t]$. The short- and long-run child penalties correspond to event time 1 and 10, respectively. The coefficients are estimated on a balanced sample of workers who had their first child between 1995 and 2009 with age at childbirth between 20 and 45. The effects are estimated unconditional on employment status, imputing zeros when the worker is not employed. Shaded areas represent 95% confidence intervals based on standard errors clustered at the individual level.

Figure A.2: Heterogeneous first-stage and reduced-form effects



Notes. The top row reports first-stage coefficients and their 95% confidence intervals from the estimation of equation (1) with remote work as the outcome. The bottom row reports reduced-form coefficients and their 95% confidence intervals from the estimation of equation (1) for the change in earnings between the year after and two years before childbirth, rescaled by mean earnings two years before childbirth. From left to right, columns report heterogeneity by mothers' age, mothers' earnings quartile, fathers' earnings quartile, and commuting status. Confidence intervals are based on standard errors clustered at the 4-digit sector and local labor market level.

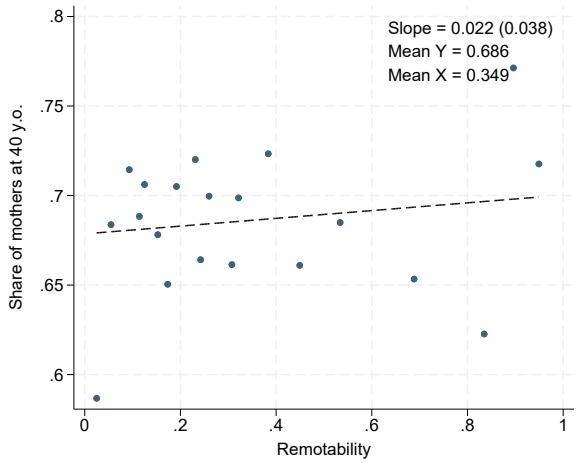
Figure A.3: First-stage dynamic estimates for fathers' WFH contracts



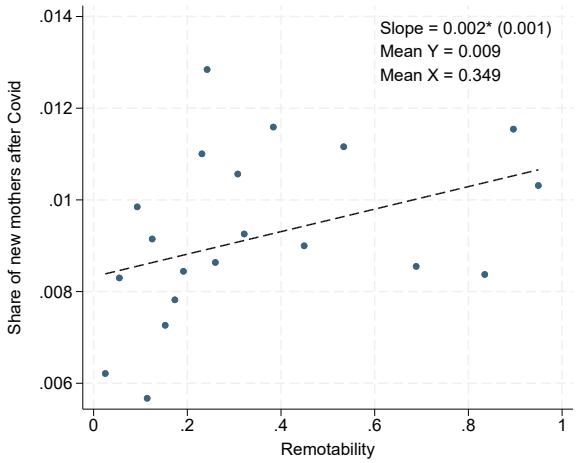
Notes. This figure plots first-stage coefficients and their 95% confidence intervals for fathers, based on standard errors clustered at the 4-digit sector and local labor market level. The two lines report estimates for the full sample of fathers and those in households where both parents are employed in the private sector at the time of childbirth. Coefficients come from estimating equation (1), where the outcome is the probability of holding a WFH contract.

Figure A.4: Sector-level correlation of fertility and remotability

(A) Share of mothers at 40 years old



(B) Share of new mothers after COVID



Notes. This figure reports binned scatter plots of the relationship between various measures of fertility and remotability at the sector level. In panel A the variable on the vertical axis is the share of mothers at the end of their fertile age for the cohorts 1974 to 1979 employed in the counterfactual exercise of Section 8. In panel B, it is the share of new mothers in a sector after 2019. Each graph reports the slope of the relationship alongside robust standard errors. Correlations are weighted by a sector's employment.

A A Framework of Child Penalty and Remote Work

A.1 Environment

A parent has a time endowment normalized to one. Childcare requirements depend on the number of children k, with $T'(k) > 0$. Time is allocated between leisure, market work, and childcare:

$$\ell_i = 1 - T_i(k) - (1 + \tau_i)h_i - \phi_i(k),$$

where $\tau_i \geq 0$ captures time costs of working (commuting, rigid schedules) and $\phi_i(k)$ captures coordination burdens from combining work and childcare that are independent of hours worked. Flexibility determines the effective magnitude of both costs. Writing $\tau(a)$ and $\phi(\cdot, a)$ with a the degree of remotability, flexible arrangements satisfy $\tau(a') < \tau(a)$ and $\phi(\cdot, a') < \phi(\cdot, a)$ for $a' > a$. Coordination costs are supermodular: $\partial^2\phi(k, a)/\partial k\, \partial a > 0$.[9]

A.2 Mother-only Model

As a benchmark, consider a single mother choosing $h \geq 0$ to maximize $U = \ln(wh) + \alpha \ln \ell$, $\alpha > 0$, subject to

$$\ell = 1 - T(k) - \phi(k, a) - (1 + \tau(a))h.$$

The optimal hours and earnings are

$$h^*(k, a) = \frac{1 - T(k) - \phi(k, a)}{(1 + \tau(a))(1 + \alpha)}, \qquad E(k, a) = wh^*(k, a).$$

The child penalty is $\Delta(k, a) \equiv E(0, a) - E(k, a)$.

Proposition 1 (Flexibility reduces the child penalty)**.** *If $\tau(a)$ and/or $\phi(k, a)$ are relevant for the mother's labor decision, greater remotability increases $E(k, a)$ and decreases $\Delta(k, a)$.*

Proof. Immediate from $h^*(k, a)$ and the fact that τ and ϕ decrease in a. □

A.3 Heterogeneous effects across mothers

A.3.1 High- vs low-earnings mothers

Let the time wedge and coordination burden depend on the mother's wage: $\tau(w, a)$ and $\phi(k, w, a)$, with $\tau_w < 0$ and $\phi_w < 0$ so that lower-earning mothers face more rigid jobs. The percentage effect of flexibility on post-birth earnings is

$$\frac{\partial \ln E(k, w, a)}{\partial a} = -\frac{\tau_a(w, a)}{1 + \tau(w, a)} - \frac{\phi_a(k, w, a)}{1 - T(k) - \phi(k, w, a)}.$$

Since lower-earning mothers have larger τ and ϕ, the denominators are smaller. With similar numerators across wage groups, both terms are larger in magnitude for low earners, implying

[9] This assumption is used only in the fertility result (Section A.5).

larger proportional gains. This formalizes the idea that larger *relative* gains for lower-earning mothers are consistent with job rigidity as a mechanism.

Larger or comparable *absolute* gains require stronger labor-supply responses to offset the mechanical effect of lower wages on $\Delta E(k, w) = w\,\Delta h(k, w)$. A natural mechanism is proximity to discrete margins. Let mothers choose among

$$h \in \{0, \bar{h}^P, \bar{h}^F\},$$

corresponding to non-employment, part-time, and full-time work. If lower-earning mothers are more likely to shift from full-time to part-time, or from part-time to non-employment, under rigid arrangements, flexibility can prevent these transitions and generate larger absolute earnings gains despite lower wages.

A.3.2 Younger vs older mothers

Let s denote career stage (lower s = younger). Future earnings depend on career accumulation:

$$E_1(s, a) = \omega(H(s) - \delta(a)), \qquad \omega' > 0,\ \omega'' < 0,$$

where $\delta(a)$ captures career disruption with $\delta_a < 0$. Total post-birth value is $V(k, s, a) = E_0(k, s, a) + \beta E_1(s, a)$.

Proposition 2 (Age heterogeneity)**.** *Suppose (i) flexibility reduces current frictions more strongly at earlier career stages: $\tau_{as} > 0$, $\phi_{as} > 0$; and (ii) flexibility mitigates career disruption: $\delta_a < 0$. Then the marginal effect of flexibility is larger for younger mothers: $\frac{\partial}{\partial s}\left(\frac{\partial V}{\partial a}\right) < 0$.*

Proof. The current earnings effect satisfies

$$\frac{\partial E_0}{\partial a} = \frac{w(s)}{1+\alpha}\left[\frac{-\phi_a}{1+\tau} + \frac{Q \cdot (-\tau_a)}{(1+\tau)^2}\right] > 0.$$

Under assumption (i), differentiating with respect to s gives $\frac{\partial}{\partial s}\left(\frac{\partial E_0}{\partial a}\right) < 0$. The future earnings effect satisfies

$$\frac{\partial E_1}{\partial a} = -\omega'\delta_a > 0, \qquad \frac{\partial}{\partial s}\left(\frac{\partial E_1}{\partial a}\right) = -\omega'' H' \delta_a < 0,$$

using $\omega'' < 0$ and $\delta_a < 0$. Summing both terms gives the result. □

A.4 Household Problem

A household consists of a mother (m) and a father (f). Under a unitary representation, spouses jointly choose labor supply and childcare allocation to maximize

$$U = \ln c + \alpha_m \ln \ell_m + \alpha_f \ln \ell_f, \qquad \alpha_m, \alpha_f > 0,$$

where $c = w_m h_m + w_f h_f$ is household consumption and leisure satisfies

$$\ell_i = 1 - \phi_i(k) - t_i - (1 + \tau_i)h_i,$$

subject to the childcare feasibility constraint $t_m + t_f = T(k)$.[10] Define the effective wage $\tilde{w}_i \equiv w_i/(1 + \tau_i)$. The first-order conditions for market hours are

$$\alpha_i \frac{c}{\ell_i} = \tilde{w}_i, \qquad i \in \{m, f\}, \tag{A.1}$$

and optimal childcare allocation requires, when both parents provide childcare,

$$\frac{\alpha_m}{\ell_m} = \frac{\alpha_f}{\ell_f}, \tag{A.2}$$

equalizing the marginal utility cost of childcare across spouses.

Why the child penalty arises for mothers. Three mechanisms in the model can generate a child penalty for mothers only, corresponding to the channels discussed in Section 2.

Comparative advantage. If $\tilde{w}_f > \tilde{w}_m$, a marginal reallocation of childcare from mother to father reduces household consumption by $\tilde{w}_f$ rather than $\tilde{w}_m$. The household therefore assigns childcare to the mother at a corner ($t_f = 0$).

Forced work interruptions. Even when effective wages are similar, some early childcare — pregnancy recovery, breastfeeding, mandatory leave — cannot be reallocated across parents. This generates a minimum disruption to mothers' labor supply. If work continuity or full-time premia matter, the household optimally avoids extending the disruption to fathers, producing an asymmetric penalty.

Gender norms. Social norms favoring maternal specialization can be represented as a higher disutility from fathers providing childcare, or as an implicit constraint on t_f.

These mechanisms are observationally similar in the absence of variation in workplace flexibility. The response to fathers' flexibility — examined in the next paragraph — helps distinguish them empirically.

Father flexibility and mothers' earnings. Greater flexibility lowers fathers' effective time wedge and coordination burden, increasing their time slack for given h_f.

Proposition 3 (Father flexibility improves mothers' outcomes)**.** *If fathers' effective time costs are relevant, greater father flexibility reallocates childcare toward fathers, relaxes mothers' time constraint, and increases mothers' hours and earnings.*

[10] The key feature of this representation is that childcare must be allocated within the household and spouses may face different opportunity costs of providing it. The qualitative predictions hold in a broad class of collective household models.

Proof. Fix $h_f = \bar{h}_f$ and let $Y_f = w_f \bar{h}_f$. Define $\bar{t}_f \equiv 1 - \phi_f(k) - (1+\tau_f)\bar{h}_f$ and $\bar{t}_m \equiv 1 - T(k) - \phi_m(k)$. The household FOCs imply $\ell_m = \frac{\alpha_m}{\alpha_f}\ell_f$, which gives

$$h_m^* = \frac{w_m(\bar{t}_f + \bar{t}_m) - (1+\tau_m)(\alpha_f + \alpha_m)Y_f}{(1+\tau_m)w_m(1+\alpha_f+\alpha_m)},$$

so $\frac{\partial h_m^*}{\partial t_f} = \frac{1}{(1+\tau_m)(1+\alpha_f+\alpha_m)} > 0$. Since $\bar{t}_f$ increases when τ_f or ϕ_f fall, the result follows. □

Employer rigidity and a zero father child penalty. The proof fixes h_f to isolate the time-reallocation channel. In the unconstrained problem, the effect of father flexibility on fathers' own hours is ambiguous: lower τ_f raises the effective wage (increasing hours), but lower ϕ_f reduces the cost of father-provided childcare (potentially reducing hours). Employer-side rigidities resolve this ambiguity.

Suppose fathers' earnings include a full-time premium $\kappa > 0$:

$$y_f(h_f) = \begin{cases} w_f h_f & \text{if } h_f < \bar{h}, \\ w_f h_f + \kappa & \text{if } h_f \geq \bar{h}. \end{cases}$$

If κ is sufficiently large, the household never reduces fathers below $\bar{h}$ after childbirth, since continuous adjustments along other margins (mothers' hours, childcare reallocation, leisure) dominate. Fathers' hours therefore remain fixed and flexibility operates exclusively through childcare reallocation. The coexistence of a null father penalty and positive father-flexibility effects on mothers' outcomes is thus consistent with a household-level time constraint interacting with rigid labor contracts, and is difficult to reconcile with mechanisms operating exclusively on the mother side.

A.5 Fertility

We embed the household problem in a simple intertemporal fertility choice. The household has k children at the start of period 1 and chooses whether to have an additional child $b \in \{0, 1\}$, entering period 2 with $k+b$ children. Let $V(k)$ denote the indirect utility from the within-period household problem at parity k. The household maximizes

$$W(b; k) = V(k) + \beta\Big(V(k+b) + B(k+b) - F(k+b)\Big),$$

where $B(\cdot)$ is the utility benefit from children and $F(\cdot)$ a fixed cost. It chooses $b = 1$ whenever

$$\Delta W(k) \equiv \beta\Big[V(k+1) - V(k) + B(k+1) - F(k+1)\Big] \geq 0. \tag{A.3}$$

Proposition 4 (Flexibility raises fertility)**.** *Suppose $V(k+1) < V(k)$, so that higher parity tightens the household time constraint. Greater workplace flexibility increases $V(k+1) - V(k)$ and therefore raises the probability of an additional birth.*

Proof. Since $B(\cdot)$ and $F(\cdot)$ do not depend on flexibility, the effect on $\Delta W(k)$ operates entirely

through $V(k+1) - V(k)$. Flexibility lowers effective time wedges and coordination burdens, expanding the feasible set at both parities. Since the constraint binds more tightly at higher parity, the relaxation raises $V(k+1)$ by more than $V(k)$ — formally guaranteed by the supermodularity of coordination costs in k and a. Hence $\Delta W(k)$ increases and the household is more likely to choose $b = 1$. □

A positive fertility response among more-exposed households is therefore consistent with the household-constraint interpretation: it implies that rigid workplace arrangements were raising the expected cost of an additional child, distorting family choices beyond post-birth earnings trajectories. This is especially informative at higher parity, where $F(\cdot)$ is smaller and the decision depends more directly on whether existing arrangements can absorb additional childcare needs — making a fertility response less likely to reflect anticipatory career concerns

B Diffusion of WFH in Italy and its determinants

WFH in Italy was rarely used before the pandemic. According to the Italian Labor Force Survey (LFS), in 2019 only 1.5% of non-farm private sector employees worked from home at least one day per week on average. The same figure spiked to 11.7% in 2021 to decline slightly to 8.5% (Figure B.1). The pandemic indeed boosted WFH adoption whose diffusion persists. WFH rose much more for women than for men (in 2024, a 6-time increase versus about 5). Our matched administrative data, despite starting in 2020, mimic such trend: they record above 2 million firm-worker pair records in 2020-2021 and about 1.8 million in 2024.

The adoption of WFH depends on several factors that vary between occupations due to the task that can be performed remotely or on site (Basso et al., 2022; Dingel and Neiman, 2020). Firms ex ante preparedness and performance during the pandemic (Basso et al., 2025), ICT investments (Basso et al., 2025; Boeri et al., 2025) and managerial practices (Lamorgese et al., 2024) have also been proven to be a determinant of WFH adoption. Personal conditions correlated with labor supply decisions, including housing arrangement and household composition, also determine whether a worker is willing to work from home (Burdett et al., 2024). To isolate the demand component of WFH, we analyze two features that we will further exploit in our empirical analysis.

Some tasks and occupations are more amenable to remote work because they do not require in-person interactions with colleagues, suppliers, or customers. Using the measure of WFH feasibility developed by Basso et al. (2022) for European labor markets, we compute sector-level WFH potential as the share of workers whose occupations can be performed remotely. Even at a granular level, this measure exhibits substantial heterogeneity across sectors. Table B.1 compares sectoral WFH potential with actual use in 2024 based on LFS data. Construction and retail display low potential, while professional and high-skill service sectors—such as information and communication, finance, and scientific services—show much higher feasibility. Actual WFH use is strongly correlated with sectoral potential.

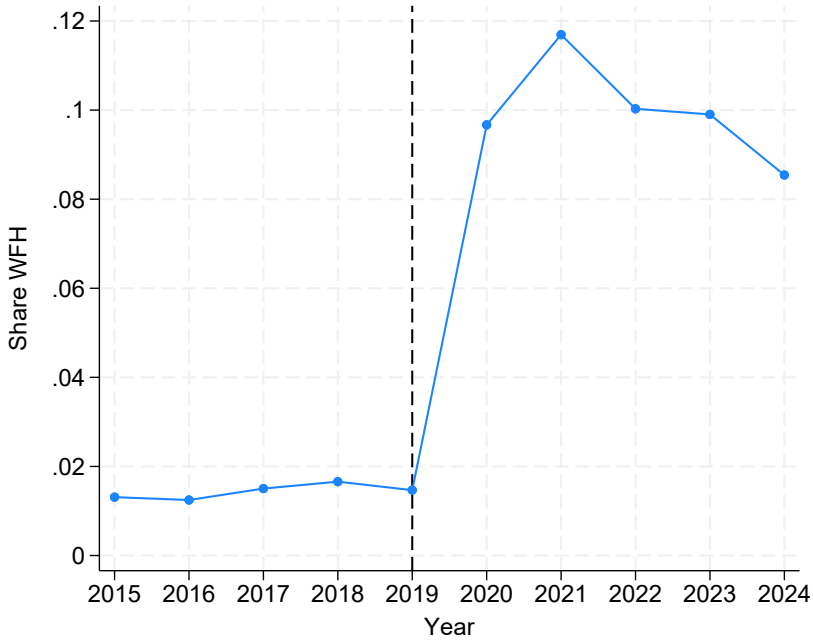
A second determinant of WFH adoption, independent of workers' labor supply decisions, is infrastructure availability. Beyond firm-level ICT investment, effective remote work requires high-speed broadband access. Using data from AGCOM, the national communications authority, we measure broadband speed at the municipal level and aggregate it to commuting zones (CZs) to capture connection quality at home and at the workplace. In 2019, broadband access varied substantially across CZs: the 90th-to-10th percentile ratio of average speed was 10.3, with a mean of 72.4 Mbps.

Table B.1: WFH: potential and 2024 use by macrosector

	(1)	(2)
	Potential WFH sh.	2024 WFH sh.
Agriculture	.041	.009
Manufacturing (food, textiles, wood)	.178	.028
Manufacturing (chemicals, pharma); water and energy	.207	.063
Manufacturing (machinery, automotive, furniture)	.235	.089
Construction	.117	.021
Wholesale and retail trade	.236	.036
Transportation; ICT services	.419	.214
Financial and professional services	.546	.187
Education and health (private)	.140	.025
Arts and other services	.128	.057

Notes. This table reports the WFH potential as of 2019 (Basso et al., 2022) and the actual use as of 2024 reported by the Italian Labour Force Survey.

Figure B.1: Trends in WFH according to the Italian Labor Force Survey



Notes. The graph displays the trends in the share of non-farm private sector employees who report to work from home at least one day in the reference week, according to the Italian Labor Force Survey.