

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

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

May 2026

Abstract

We study whether workplace flexibility is a key driver of the motherhood penalty. Exploiting the sharp and heterogeneous diffusion of work-from-home (WFH) contracts after COVID-19 in a difference-in-differences framework, we combine administrative data on the universe of Italian WFH contracts with matched employer–employee records, mother–father links, fertility records, and firm balance sheets. Greater exposure to flexible work substantially reduces mothers’ post-childbirth earnings losses, primarily through higher weeks worked, lower part-time incidence, lower parental leave take-up, and improved career progression. IV estimates indicate that holding a WFH contract offsets about 77% of mothers’ earnings losses around childbirth. Gains are larger among younger, lower-earning, and commuting mothers, in households where the father is the dominant earner, and in occupations with steeper hours–earnings profiles, consistent with flexibility relaxing constraints where most binding. Fathers’ own earnings do not respond to flexibility around childbirth, yet fathers’ exposure to flexible work reduces mothers’ earnings losses by a comparable magnitude, pointing to household-level time constraints as a central mechanism. Consistent with this interpretation, fertility, a joint household decision, rises for women more exposed to flexibility, as 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 10.7% smaller under current WFH diffusion and up to 29.4% smaller if all remotable jobs had adopted flexible arrangements.

Keywords: Work-from-home; Motherhood 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 *motherhood 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 motherhood 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 motherhood 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 motherhood 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 motherhood penalty by exploiting variation in exposure to remote work across sectors and locations. Beyond estimating the impact of a specific work arrangement, which we interpret as a measurable margin of broader workplace flexibility, we use WFH diffusion to assess whether rigid workplace organization is a central driver of the earnings divergence following childbirth. In doing so, the paper connects two strands of the literature—on the sources of the motherhood penalty and 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 exploiting the sharp post-COVID diffusion of WFH (e.g., [Barrero et al., 2023](#); [Basso et al., 2025](#); [Lee, 2023](#); [Monte et al., 2026](#)). We estimate a difference-in-differences design on repeated cross-sections of mothers, comparing labor outcomes across sector–locations with different pre-COVID exposure to remotable jobs. Exposure is measured as the interaction between the sector-level share of jobs feasible to perform remotely and local pre-pandemic broadband coverage. Identification relies on parallel trends between more and less exposed sector–locations, supported by the absence of pre-pandemic differential trends in mothers’ labor outcomes. Since we focus on the medium-run response and find no effect of exposure on firm performance, contemporaneous COVID-related shocks are unlikely to confound our results. A remaining concern is that the composition of mothers may change over time as WFH affects fertility and career choices. We address this by using pre-birth earnings to control for time-invariant unobservables, netting out selection effects and isolating the reduction in the motherhood penalty attributable to flexible work. Placebo exercises further show no changes in selection based on pre-childbirth earnings growth.

We begin by documenting the reduced-form effects of exposure to remotable work on mothers’ labor market outcomes. Moving from the 10th to the 90th percentile of exposure—roughly comparing a manufacturing job to a web-related one—raises mothers’ post-childbirth earnings by about 980 euros, along with 0.53 more weeks worked, higher full-year employment, and declines in part-time work and parental leave. We find no effect on employment probability or wages, indicating that flexibility primarily expands labor supply along the intensive margin. We find some evidence of changes in selection: mothers in more exposed sector–locations have higher pre-birth earnings after 2020. Netting out these pre-birth differences, post-childbirth outcomes remain positive and significant, implying that flexible arrangements attenuate earnings losses rather than reshuffling who becomes a mother. Placebo exercises reinforce this interpretation along two dimensions. First, assigning childbirth two years earlier yields no meaningful effects, indicating that mothers in more exposed sector–locations do not exhibit systematically faster earnings growth prior to childbirth. Second, assigning childbirth three years later yields no additional gains. The effects arise around childbirth, when time constraints are most binding, persist without further increases, and are not driven by confounders.

Next, we translate the reduced-form estimates into the causal effect of holding a WFH contract, instrumenting take-up with sector–location exposure. We interpret a WFH contract as capturing a bundle of changes in work organization, including greater schedule autonomy, lower commuting costs, and more family-friendly practices. The first stage is strong: a 10th–to–90th percentile increase in exposure raises the likelihood of a WFH contract by 13 percentage points—a meaningful shift, as about 34% of remote workers

work from home for at least half their time. Holding such a contract reduces the average earnings drop around childbirth by 77%, or 24% of pre-birth earnings. This large effect is consistent with the evidence above: about 90% of the earnings drop reflects intensive margin adjustments, the margin through which we show flexibility operates. We then restate the same gains in motherhood-penalty units, that is, as percentage-point reductions in the percentage earnings drop around childbirth. We use pre-birth earnings as a proxy for earnings potential and adjust for selection driven by time-invariant unobservables. WFH contracts reduce the short-run motherhood penalty by 29–34 percentage points, or 40%–67% depending on the calibration. Rigid work arrangements therefore account for a substantial share of the post-childbirth earnings divergence.

To further support the exclusion restriction, we show that sector–location exposure has no significant effect on employers’ outcomes—size, sales, and value added—or on coworkers’ earnings around childbirth, ruling out changes in firm-wide rent-sharing and broader shocks as confounders. Combined with the absence of effects outside the childbirth window, this confirms that the mechanism operates through mothers’ contractual arrangements specifically around childbirth, rather than through general improvements in firm performance or local labor market conditions.

We then study heterogeneity across subgroups of mothers. Flexible arrangements generate the largest gains where time and organizational constraints are most binding. Lower-earning and younger mothers benefit disproportionately, as do non-college mothers. Household structure also matters: mothers benefit more when fathers have lower earnings and, even more sharply, when the father is the dominant earner before childbirth, consistent with flexibility relaxing household-level constraints in more traditional earnings structures. Finally, gains are larger for mothers who commuted before childbirth and for those in greedy occupations, where earnings are more strongly tied to long, inflexible hours. Overall, these results show that flexible work matters most where rigid arrangements are most costly around childbirth, highlighting the role of household-level constraints in shaping the labor cost of motherhood.

We also examine fathers’ responses and the role of household-level constraints. Consistent with the limited fatherhood penalty, fathers’ own exposure has little effect on their earnings or labor supply around childbirth. Yet fathers’ exposure independently reduces mothers’ earnings drop by up to about 630 euros—comparable to the effect of mothers’ own exposure—suggesting that fathers’ flexibility matters not by increasing their labor supply but by relaxing household time constraints. This partner effect is identified conditional on mothers’ exposure, absent in placebo specifications away from childbirth, and stronger when mothers work in less remotable jobs. These results indicate that the benefits of flexibility arise precisely when childcare needs become binding and operate through intra-household reallocation: the motherhood penalty is not only a mother-specific labor market phenomenon, but reflects joint household constraints interacting with workplace

rigidity for *both* spouses.

Consistent with workplace rigidity operating as a household-level constraint, we examine whether flexible arrangements affect fertility decisions. If the motherhood penalty reflects the labor-market cost of children, greater flexibility lowers the marginal cost of an additional child. Following two cohorts of working women over four years—one before (placebo) and one after the COVID shock—we find an increase in births in more exposed sector-locations, on both the intensive margin (1.3% of the placebo average, 15% in the IV) and the extensive margin of entry into motherhood (0.44%, 3% in the IV). Effects are smaller on the latter, consistent with flexibility only partially offsetting the fixed costs of a first child. A back-of-the-envelope calculation suggests that the observed WFH diffusion raises the total fertility rate by 1.2%, a figure that would increase to 3.9% if all remotable jobs adopted WFH contracts.

Finally, we assess the implications of flexible work for the lifecycle gender earnings gap. Focusing on cohorts born between 1974 and 1979—who completed their fertile years before COVID-19—the female-to-male earnings ratio falls by 19 percentage points between ages 25 and 40. Removing the full motherhood penalty leaves the ratio flat or slightly increasing, confirming that the penalty explains nearly the entire widening (Landais et al., 2025). Under the current WFH adoption, the drop would shrink by 2.0 percentage points (10.7% of the observed change); if all remotable jobs offered WFH contracts, by 5.6 percentage points (29.4%). Although WFH cannot close the entire gender earnings gap, it can meaningfully slow its lifecycle growth.

We contribute to two strands of the literature. First, we contribute to the literature on the determinants of the motherhood penalty (Angelov et al., 2016; Kleven et al., 2024a,b, 2019), which has documented large and persistent gender gaps arising at childbirth but provided limited causal evidence on their drivers. A prominent explanation emphasizes workplace organization and the lack of temporal flexibility (Albanesi, 2026; Goldin, 2014, 2022), alongside institutional and household factors shaping labor supply and specialization (Cortés and Pan, 2023; Olivetti and Petrongolo, 2016, 2017). These explanations reflect a central theme in the recent literature: the distinction between preferences and constraints in shaping gender gaps. Much of the remaining inequality in advanced economies is understood to reflect differential work-family constraints rather than differences in preferences or skills (Olivetti et al., 2024), making workplace flexibility central to the debate. However, most existing evidence relies on survey data or occupation-based proxies, making it difficult to isolate the causal role of contractual arrangements. Pre-pandemic US studies show that WFH feasibility reduces the motherhood employment penalty but may induce sorting into lower-paying firms (Harrington and Kahn, 2025; Jack et al., 2025); European and Latin American evidence associates remote-work feasibility with narrower gender gaps (Arntz et al., 2022; Crescenzi et al., 2025; Zarate, 2025); pandemic-era evidence is more mixed, likely reflecting contemporaneous child-

care disruptions (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 actual WFH contract adoption and exploit their rapid post-COVID diffusion as plausibly exogenous variation in workplace flexibility, allowing us to trace effects on mothers’ earnings, employment, and career progression, examine cross-parent spillovers, study fertility decisions, and—most importantly—provide direct causal evidence that rigid workplace organization is a central driver of the motherhood 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 single firms—show that remote and hybrid work affect retention, satisfaction, and career progression (Atkin et al., 2023; Barrero et al., 2023; Bloom et al., 2024, 2015; Choudhury et al., 2024; Mas and Pallais, 2017), while 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 show that the effects of WFH are highly state-dependent, concentrated around parenthood rather than reflecting a uniform premium or penalty, and shaped by household constraints including fathers’ flexibility. Our analysis is closely related to Aksoy et al. (2026) and Lu et al. (2025), who document a positive association between WFH and fertility using survey data. We complement their results with causal evidence, exploiting the sudden expansion of formal WFH contracts in an individual-level difference-in-differences design on population-wide administrative data. By studying fertility jointly with post-birth labor supply and earnings within couples, we offer a more comprehensive assessment of how rigid work arrangements shape household behavior and contribute to the motherhood penalty.

2 Conceptual Framework

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

Flexibility and the Motherhood 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 motherhood 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 effects of flexibility serves as a diagnostic for the sources of the motherhood penalty. Along mothers' own characteristics, larger effects for younger, lower-earning, or non-college mothers would indicate that rigidity is most costly where labor market attachment is more fragile and adjustment occurs along discrete margins such as part-time work or non-employment. Household characteristics speak to a different mechanism. If the penalty reflects a joint time constraint within the household, flexibility should be more valuable in two cases. First, it should matter more when fathers have lower earnings, either because the household has less capacity to absorb childcare costs or because the opportunity cost of shifting some childcare to the father is lower. Second, it should matter more in households where fathers are the dominant earners, as in these cases mothers are more likely to bear the labor-supply adjustment after childbirth. Finally, job characteristics provide a direct test of workplace rigidity: effects should be larger for commuters, for whom flexibility saves time directly, and for mothers in greedy occupations, where earnings are more strongly tied to long and inflexible hours. Together, these dimensions allow us to distinguish mother-specific labor market frictions from household-level constraints and job-specific sources of rigidity.

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 motherhood 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.

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 than on the fixed costs of entering parenthood.

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 2012 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, a broad occupational tier 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.

Childbirth 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 and father in

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

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% for children born between 2016 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. Since our main results on mothers’ outcomes are qualitatively and quantitatively similar in the two samples, we do not expect the household-level results to be driven by the slight overrepresentation of lower-income households in the UCA data.

Education and Detailed Occupation. We supplement Social Security records with information from the *Comunicazioni Obbligatorie*, the administrative system collecting mandatory notifications that employers submit to the Italian Ministry of Labour when employment contracts are activated, transformed, or terminated. These data are available at the contract-spell level from 2010 onward and report, among other variables, workers’ education and detailed occupation. Since the data are based on employment flows, they do not contain information for all workers in our analysis sample.² For each worker, we define college education as an indicator equal to one if the worker ever reports a tertiary degree in the available records. For occupation, we use the last observed occupation before childbirth in the firm where the woman is employed at childbirth. Occupation is recorded using the Italian occupational classification maintained by the National Statistical Institute (Istat), which we harmonize to ISCO-08.

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

²We retrieve occupation information for approximately 77% of workers in the analysis sample and education information for approximately 93%.

information on maternity episodes happening before our observation window. We use these data—whose sample size is of 81,226 mothers³—to study the implications of WFH 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 WFH contracts in Italy, drawn from mandatory “lavoro agile” declarations introduced in 2017. Defined by law as a form of subordinate employment without time or space constraints, it therefore corresponds to WFH. “Lavoro agile” was introduced to help workers balance work and family life while fostering productivity. Firms are required to report WFH spells at the worker-week level through an online portal, as workers are entitled to public insurance coverage for accidents and occupational diseases even outside the company premises. Records are consistently available since 2020 and include the firm identifier, employee gender, year of birth, citizenship, and spell start and end dates, allowing us to match these data to social security records.

We aggregate the data to the year level and create firm-level cells based on gender-age-citizenship combinations. When social security records contain more employees than WFH contracts in a given cell, we assign each employee a probability of holding a WFH contract equal to the ratio of WFH contracts to employees in the cell. In 15% of cases we find one-to-one matches; in 77% of cases there are more employees than contracts, though in 11% of these the discrepancy is at most two. In the 8% of cases where contracts exceed employees, we winsorize the probability to one. Since we instrument this probability, any measurement error in WFH take-up is not a concern for our estimates.

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%), and earn less (16,531 vs. 20,192 euros). They are also less likely to hold managerial positions and are more concentrated in

³After applying the sample restrictions—namely, restricting the sample to mothers aged 25 to 40 and with at least three years of private-sector work experience before age 30—we are left with 39,232 mothers.

Table 1: Characteristics of mothers and fathers

	(1)	(2)	(3)	(4)
	All		Households	
	Mothers	Fathers	Mothers	Fathers
Age	32.536 (4.913)	34.309 (6.156)	31.918 (4.778)	34.525 (5.609)
Full-time	0.589 (0.492)	0.831 (0.374)	0.620 (0.485)	0.865 (0.342)
Permanent	0.863 (0.344)	0.811 (0.392)	0.862 (0.345)	0.856 (0.351)
Manager	0.020 (0.141)	0.033 (0.179)	0.020 (0.141)	0.043 (0.202)
Weeks worked	34.956 (18.853)	37.101 (19.931)	36.365 (18.439)	40.533 (18.168)
Earnings	16,530.982 (13,024.498)	20,192.350 (18,038.508)	17,230.613 (12,695.311)	22,975.725 (18,458.016)
Prob. WFH	0.064 (0.222)	0.058 (0.205)	0.076 (0.240)	0.073 (0.229)
South	0.252 (0.434)	0.289 (0.453)	0.224 (0.417)	0.223 (0.416)
Industry	0.193 (0.395)	0.420 (0.494)	0.199 (0.399)	0.423 (0.494)
Firm size 1–14	0.390 (0.488)	0.315 (0.465)	0.362 (0.481)	0.287 (0.453)
Firm size 15–99	0.240 (0.427)	0.295 (0.456)	0.246 (0.430)	0.291 (0.454)
Firm size 100+	0.370 (0.483)	0.390 (0.488)	0.392 (0.488)	0.421 (0.494)
Coworker earnings	18,621.186 (14,103.524)	21,737.113 (14,602.819)	19,207.725 (14,250.261)	23,838.791 (15,150.148)
Male coworker earnings	22,549.172 (17,510.879)	22,954.246 (16,205.376)	23,207.078 (17,573.365)	25,316.705 (16,938.035)
Value added per worker	46.144 (51.649)	57.320 (55.061)	47.724 (52.567)	62.121 (57.308)
Sales per worker	204.695 (338.566)	236.206 (333.992)	211.845 (344.697)	261.419 (351.550)
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.

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 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, according to Labor Force Survey data, 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% (see Appendix C and Figure B.1). The increase was more pronounced for women than for men. Among employees who work from home, about 34% do so for at least half of their working time in 2024. This intensive use of remote work is broadly similar among men and women.

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, constructed using the Italian Labor Force Survey and defined as the pre-pandemic share of occupations that can be performed remotely based on their 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.⁴

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

⁴Beyond 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 (*Sistemi locali del lavoro*), 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.

pandemic hit.⁵ 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 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 baseline 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

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

where i indexes a mother, t is the calendar time at which she gives birth, α_s is a sector×location fixed-effect, λ_t is a calendar year fixed-effect, $\text{Remotable}_{s(i)}$ is the remotability of the job held by i at childbirth, and $X_{\bar{s}(i)}$ is a two-digit sector fixed-effect that allows us to control non-parametrically for heterogeneous time trends across sectors. 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, weeks worked and weeks of parental leave use, contract type, and promotion to managerial jobs as 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 interpret the magnitudes, in the reduced-form analysis we rescale these coefficients to reflect a move from the 10th to the 90th percentile of the $\text{Remotable}_{s(i)}$ distribution

⁵Personal 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.

among working mothers. This corresponds, holding broadband fixed, to shifting from low remotable sectors (*e.g.*, aircraft manufacturing) to high-remotability ones (*e.g.*, web portal activities), and, holding remotability fixed at a high level, to moving from low-speed areas (*e.g.*, Brindisi) to high-speed ones (*e.g.*, Milan).

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

$$Y_{it} = \alpha_{s(i)} + \lambda_t + \beta_{\text{Post}} \cdot \text{Post}_t \cdot \text{WFH}_i + \sum_{k=2016}^{2024} \gamma_k \cdot I(t = k) \cdot X_{\bar{s}(i)} + \varepsilon_{it}, \quad (2)$$

where we instrument the likelihood of holding a WFH contract (WFH_i) with the sector-location exposure to remotable work ($\text{Remotable}_{s(i)}$), and Post_t is a dummy for the post-2019 period that we include both in first-stage and reduced-form regressions.

We interpret WFH as a measurable and economically important margin of workplace flexibility, while recognizing that its adoption may coincide with broader changes in work organization—such as greater schedule autonomy, lower commuting costs, or more family-friendly practices. We therefore view our estimates as capturing the effect of a bundle of flexibility-improving changes centered on WFH, rather than the effect of performing work remotely in isolation.

Identification Assumptions and Exclusion Restriction. Our identification relies on the parallel trends assumption: in the absence of the pandemic, outcomes for mothers in sector-locations with different levels of remotability would have followed similar paths. We consider two categories of potential confounders.

The first is time-varying shocks at the sector-location level that may be correlated with $\text{Remotable}_{s(i)}$. These include both differential pre-trends and demand shocks occurring simultaneously to the pandemics. We address the former by showing that pre-2020 coefficients from equation (1) are flat, providing direct support for parallel trends. As for demand shocks, these are less of a concern because we focus on the medium run—by which point COVID-induced shifts had largely reversed⁶—and because our specifications include sector-by-time fixed effects that absorb persistent sectoral changes. Our results are also robust to including location-time fixed effects, which might capture heterogeneous changes in local amenities unrelated to contract flexibility. Importantly, if sector-location demand shocks were driving the results, we would expect to see effects on firms or coworkers as well. In Section 5.2, we find no effects on firm outcomes (size, sales, value added) or coworkers’ earnings. This evidence helps rule out sector-location shocks as a confounding factor and supports the exclusion restriction for our IV estimates, which requires that remotability affects mothers’ outcomes only through work arrangements rather than through changes in firm productivity or local labor market conditions.

⁶COVID-19 generated large sectoral demand shifts (Coibion et al., 2025; Guerrieri et al., 2022), which largely reversed as restrictions eased.

The second category of confounders is individual-level time-varying shocks. Since we estimate our model on a repeated cross-section of mothers, we require that individual unobservables are uncorrelated with $\text{Remotable}_{s(i)}$. This assumption is violated if WFH availability changed the composition of mothers over time. Suggestive evidence on potential selection can be obtained by examining whether pre-childbirth employment histories and observable characteristics of mothers vary with our instrument. We also address this concern directly by using pre-birth earnings as a proxy for mothers' earnings potential, which allows us to net out compositional changes driven by time-invariant unobservables and isolate the reduction in the motherhood penalty attributable to flexible work arrangements. While using first differences accounts for time-invariant unobservables, estimates may still be biased by changes in the selection of mothers based on their earnings growth. To address this concern, in Section 5.1 we conduct placebo tests showing that mothers in more exposed sector-locations do not exhibit faster earnings growth prior to childbirth.

Quantifying the Motherhood Penalty Reduction. We formalize the strategy that exploits pre-birth earnings using a simple decomposition of post-childbirth earnings. Let Y^c be post-childbirth earnings and $\tilde{Y}$ mothers' earnings potential, with $Y^c = (1 - \rho)\tilde{Y}$, where ρ summarizes the post-childbirth earnings penalty. Changes in observed post-childbirth earnings can therefore reflect either a change in the motherhood penalty or a change in the composition of mothers. Our baseline outcome, the change in earnings around childbirth, already removes level differences in pre-birth earnings and measures the reduction in the penalty in euros.

We next express the same effect as a percentage reduction in the motherhood penalty, proxying earnings potential with pre-birth earnings. Exploiting the relationship between Y^c and $\tilde{Y}$, we implement


$$d\rho = \frac{\beta^{Y^c} - (1 - \rho)\beta^{\tilde{Y}}}{\tilde{Y}}, \quad (3)$$

where β^Y is the difference-in-differences estimate of the effect of WFH on a generic outcome Y . This requires estimates of WFH on pre- and post-childbirth earnings, a calibrated value for the baseline motherhood penalty, and average pre-birth earnings.

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

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.

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 Arrangements and Mother's Outcomes

5.1 Effects on Earnings, Employment, and Promotions

We begin by estimating specification (1) on mothers' outcomes in the years following childbirth. Figure 2, Panel A, plots the estimated coefficients and Table 2 summarizes the results for all the outcomes discussed in this Section, reporting estimates of a static difference-in-differences regression (i.e., replacing year interaction terms in equation (1) with an indicator for the years after 2019). We report the effect of moving from the 10th to the 90th percentile of the exposure distribution in the sample of mothers. We find that greater availability of flexible work arrangements has a positive and increasing effect on mothers' earnings in the year after childbirth. Between 2021 and 2023, post-childbirth earnings are 1,100–1,300 euros higher in more exposed sector–location cells, consistent with remote work relaxing the time and coordination constraints that bind most tightly around childbirth. The effect phases in gradually: since COVID-19 struck in early 2020, only mothers giving birth in the final months of that year were exposed to

the new flexibility regime, with full exposure only from 2021 onward.

As discussed in Section 4, part of the increase in earnings may reflect changes in the composition of mothers. To address this, we subtract pre-birth earnings from the outcome to account for shifts in mothers' earnings potential.⁷ When we use the change in earnings around childbirth as the outcome, we continue to find positive effects, though slightly smaller: the estimated earnings change is about 410 euro in 2021, and just above 1,000 euro in 2023, confirming that WFH opportunities substantially mitigate the motherhood penalty. The smaller magnitude relative to post-birth earnings alone suggests that pre-birth earnings also increased in more exposed sector–locations, consistent with a detectable but modest compositional shift toward higher-earning mothers after the pandemic. The beneficial effect of flexibility is not short-lived. When we consider the change in earnings between two years after and two years before childbirth, the estimates remain positive and grow over time, reaching about 590 euro in 2021 and 740 euro in 2022, compared to 646 euro when we measure post-childbirth earnings one year after birth for the same 2022 cohort. This suggests that the reduction in the motherhood penalty persists beyond the first year after birth, consistent with flexible work having a lasting effect on mothers' labor market attachment.


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. Netting out the change in weeks worked before childbirth, the effect remains positive and larger—around 1.7 additional weeks in 2023—indicating that compositional changes do not drive the observed increase in labor supply. In Figure A.1, we show that weekly earnings conditional on working do not change, and that, conditional on working, the percentage increases in total earnings and weeks worked are nearly identical. This implies that flexible work operates mainly through the intensive labor supply margin by expanding the feasible allocation of time between market work and childcare, with little role for wage increases. The increase in weeks worked is further mirrored by a decline in parental leave of about 0.5–0.7 weeks after 2020, bringing additional evidence that mothers were facing a binding time constraint.

Consistent with the intensive-margin results, the probability of employment is largely unaffected (Panel C). Most mothers remain attached to the labor force regardless of flexible work availability. However, we uncover other extensive margin responses. The likelihood of being employed full-year is 2.1 percentage points higher in 2023, equivalent to a 6.4% increase relative to the pre-COVID average.⁸ Moreover, the probability of holding

⁷We use earnings measured two years before childbirth as the reference period rather than the year immediately preceding birth, as the latter may already reflect maternity leave or pregnancy-related absences.

⁸Following Blau and Kahn (2017), we define a full-year employee as an individual working at least 36 full-time-equivalent weeks (three quarters). This accounts for the institutional setting: mothers are entitled to mandatory maternity leave (up to five months) and voluntary parental leave (up to six months per parent), with an average duration of approximately one quarter in 2023. Our definition therefore

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



Notes. This figure reports reduced-form coefficients and their 95% confidence intervals from the estimation of equation (1) for a range of labor market outcomes. Confidence intervals are based on standard errors clustered at the 4-digit sector and local labor market level. Panel A considers earnings outcomes: annual earnings in the year after childbirth, the change in annual earnings between one year after and two years before childbirth, and the corresponding change between two years after and two years before childbirth. Panel B considers labor supply outcomes: total full-time-equivalent weeks worked in the year after childbirth, the change in full-time-equivalent weeks between one year after and two years before childbirth, and weeks of parental leave. Panel C considers employment margins: an indicator for being employed, an indicator for being employed for at least 36 full-time-equivalent weeks, and an indicator for working part-time conditional on employment. Panel D considers a career outcome: an indicator for promotion to a managerial position in the year after childbirth, conditional on employment.

a part-time contract decreases by 1.4–2.3 percentage points for more exposed mothers (3.1–5.1% of the pre-2020 average), suggesting substitution toward full-time work when flexible arrangements are available. This substitution is particularly informative: in the absence of flexibility, part-time arrangements may function as a second-best adjustment

accounts for mothers who may be on leave for part of the year while remaining strongly attached to the labor market.

to rigid full-time norms. When remote work becomes available, mothers 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. Panel D shows that promotions to managerial positions in the year after childbirth increase by 0.3–0.5 percentage points, approximately three times the pre-2020 promotion rate among new mothers.⁹ This suggests that flexibility not only affects 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 career continuity and lower the probability that motherhood leads to persistent downward adjustments in occupational status.

Placebo Events and Timing of Childbirth. To assess whether time-varying confounders might drive our results, we conduct a set of placebo exercises using the same sample of mothers. In addition, we exploit this exercise to verify that the estimated effects are specific to the period around childbirth, when time constraints are most binding.

First, we assign to each mother a placebo childbirth year occurring two years before the actual event and re-estimate our baseline specification. We find no significant effects on changes in earnings (Figure A.2). This placebo serves two purposes. It confirms that the gains we document are concentrated around the childbirth event, and it alleviates concerns that mothers more exposed to remotability after the COVID shock were on systematically steeper earnings trajectories. Indeed, this exercise shows that these mothers did not exhibit faster earnings growth prior to childbirth.¹⁰

Second, in Figure A.2 we perform a complementary placebo by assigning a childbirth year three years after the actual one. Because the motherhood penalty emerges around childbirth and stabilizes within a few years, we do not expect additional gains from flexibility beyond those already realized if the effects are concentrated around the childbirth period. Consistent with this prediction, we find no meaningful additional effects on earnings changes in this exercise.

Additional Robustness. Our results are robust to including province-by-year fixed effects, which absorb time-varying local shocks such as changes in childcare availabil-

⁹Although sizable in relative terms, this effect is not large enough to generate meaningful changes in average weekly earnings, as shown in Figure A.1.

¹⁰Because we showed that pre-birth earnings levels are affected by changes in selection after the pandemic, level outcomes for the placebo exercise mechanically reflect patterns documented in the main analysis. For this reason, we focus here on outcomes expressed as changes around the placebo childbirth year.

ity, local amenities, and other place-specific post-COVID conditions that could act as confounders and violate the exclusion restriction. The estimates are slightly smaller, as expected, because this specification also absorbs the broadband-driven component of the exposure measure (Table A.1, panel B). However, the effects remain sizable and statistically significant on all outcomes. This indicates that the main results are not driven by heterogeneous local shocks: even after absorbing such variation, exposure to more flexible jobs improves post-birth outcomes for mothers.

Because the identifying variation is driven by a common aggregate shock (COVID), residuals may be correlated across units within a given year. Clustering at the year level provides a conservative check for such cross-sectional dependence. Given the limited number of clusters, we implement a wild bootstrap (Table A.1, panel C). The resulting inference remains unchanged.

Quantifying the Effects of Work-From-Home Contracts. We next quantify the causal effect of holding a WFH contract by instrumenting its probability with our exposure measure, estimating equation (2). Beyond the impact of a specific work arrangement, this strategy provides suggestive evidence on how much of the motherhood penalty can be attributed to rigid job structures. Table 2 reports reduced-form and 2SLS results in Panels A and B. The first stage is strong: moving from the 10th to the 90th percentile of Remotable_s raises the probability of a WFH contract by 13 percentage points (KP F-statistic: 360). The implied exposure is also economically meaningful, as in survey data about 34% of remote workers work from home for at least half their working time.

The IV estimates point to sizable effects. Holding a WFH contract increases mothers' post-childbirth earnings by almost 7,500 euro. Focusing on the change in earnings around childbirth as the outcome, the effect amounts to about 3,900 euro, relative to an average drop of 5,100 euro and pre-birth average earnings of 16,500 euro. This implies that a WFH contract reduces the post-birth earnings drop by 77%, or 24% of pre-birth earnings.

The IV identifies the effect of flexibility for mothers at the margin of adopting WFH arrangements in response to sectoral exposure. In our setting, WFH might capture a bundle of flexibility adjustments, including both where work is performed and how work time is organized. However, it does not capture effects for inherently non-remotable jobs, nor does it eliminate other sources of rigidity. The estimate thus reflects a partial relaxation of contractual constraints within existing employment relationships. From this perspective, the magnitude is particularly striking: if this single bundle of flexibility adjustments offsets most of the observed post-birth earnings decline, rigid job structures account for a substantial portion of the motherhood penalty, at least for the jobs that can be performed from home.

Finally, the magnitude of the effect (77% of the average earnings drop around childbirth) is large but consistent with the margin through which flexibility operates. In our

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

	(1) Earnings	(2) Net earnings	(3) Net earnings $t + 2$	(4) Weeks	(5) Net weeks	(6) PL weeks	(7) Employed	(8) Employed full-year	(9) Part-time	(10) Promoted post
<i>Panel A: Reduced form</i>										
Remotable $\times$ Post	978.57*** (149.84)	512.86*** (105.81)	501.45*** (110.20)	0.53*** (0.14)	0.70*** (0.16)	-0.71*** (0.10)	-0.001 (0.002)	0.015*** (0.004)	-0.022*** (0.004)	0.003*** (0.001)
Dep. var. mean	11395	-5105	-3257	25.6	-9.39	11.6	.86	.329	.457	.00192
Observations	770,625	770,625	679,226	770,625	770,625	770,625	770,625	770,625	661,562	661,562
<i>Panel B: 2SLS</i>										
WFH $\times$ Post	7502.35*** (1106.84)	3931.92*** (826.34)	3905.43*** (909.99)	4.05*** (1.02)	5.36*** (1.18)	-5.46*** (0.71)	-0.008 (0.016)	0.118*** (0.029)	-0.162*** (0.027)	0.023*** (0.004)
First stage	0.130*** (0.007)	0.130*** (0.007)	0.130*** (0.007)	0.130*** (0.007)	0.130*** (0.007)	0.130*** (0.007)	0.130*** (0.007)	0.130*** (0.007)	0.130*** (0.007)	0.130*** (0.007)
KP F-stat	360.361	360.361	331.545	360.361	360.361	360.361	360.361	360.361	363.513	363.513
Dep. var. mean	11395	-5105	-3257	25.6	-9.39	11.6	.86	.329	.457	.00192
Observations	770,625	770,625	679,226	770,625	770,625	770,625	770,625	770,625	661,562	661,562

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 are: annual earnings in the year after childbirth (column 1); the change in annual earnings between one year after and two years before childbirth (column 2); the change in annual earnings between two years after and two years before childbirth (column 3); total full-time-equivalent weeks worked in the year after childbirth (column 4); the change in full-time-equivalent weeks between one year after and two years before childbirth (column 5); parental leave duration in weeks over the two years after childbirth (column 6); an indicator for being employed in the year after childbirth (column 7); an indicator for being employed for at least 36 full-time-equivalent weeks in the year after childbirth (column 8); an indicator for having a part-time contract, conditional on employment, in the year after childbirth (column 9); and an indicator for being promoted to a managerial position, conditional on employment, in the year after childbirth (column 10). Standard errors, clustered at the 4-digit sector and local labor market level, are reported in parentheses. The first-stage coefficient and the Kleibergen-Paap F-statistic are reported in Panel B. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

sample of mothers, the extensive margin accounts for only a small share of the post-birth earnings decline: the drop in employment, scaled by average pre-birth earnings conditional on employment, explains about 10% of the total loss. The remaining 90% reflects changes within employment, including fewer weeks worked, reduced hours, and slower career progression. WFH operates precisely along this quantitatively dominant margin by relaxing constraints within ongoing employment relationships. The IV effect is therefore large but not disproportionate: flexibility mitigates the component of the motherhood penalty that accounts for most of the observed earnings loss.¹¹

Estimating the Reduction in the Motherhood Penalty. We next translate the IV estimates into an implied change in the motherhood penalty using equation (3), taking pre-birth earnings as a proxy for earnings potential. The key parameter, ρ , is calibrated in three ways. First, using our administrative data and following Kleven et al. (2019), we construct a balanced panel of first-time mothers and fathers around childbirth and estimate a first-year post-birth earnings penalty of $\rho = 0.460$ (Figure A.3, panel A). Second, following Casarico and Lattanzio (2023), we match each mother at first childbirth to a childless woman with similar observable characteristics, construct a balanced panel of earnings normalized by pre-birth earnings, and estimate a first-year penalty of $\rho = 0.496$ (Figure A.3, panel B). Third, as a robustness check, we compute ρ as the ratio of the change in earnings between the pre- and post-childbirth periods to pre-birth earnings, which yields $\rho = 0.370$. In all cases, we use average pre-birth earnings of women in the post-COVID period as the denominator in equation (3).


Across all calibrations, WFH contracts imply a sizable reduction in the motherhood penalty: 33.2 percentage points, or 61% of the full penalty, when ρ is calibrated using the mother–father gap; 34.1 percentage points, or 68%, when ρ is calibrated using the mother–placebo non-mother gap; and 29.2 percentage points, or 40%, when ρ is calibrated using the post-birth earnings decline.

5.2 Exclusion Restriction: Effects on Firms and Coworkers

For our IV estimates to be causal, the exclusion restriction requires that exposure to flexible arrangements affects mothers’ outcomes only through their contractual agreement. Our placebo exercises already suggest that selection into motherhood based on earnings growth is unlikely to confound the estimates. A remaining concern is that flexible arrangements may directly affect firm productivity, thereby influencing earnings growth through channels other than the contract itself.

¹¹The implied increase in weeks worked is mirrored by a reduction in parental leave, which is typically compensated at about 30% of wages. As a result, the estimated increase in earnings—net of these benefits—should be viewed as an upper bound on the true income gain. However, the persistence of the effect at a similar magnitude two years after childbirth, when parental leave is no longer available,

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.

To examine this, we estimate equation (1) on changes in firm- and coworker-level outcomes between the year following and preceding childbirth, with results in Figure 3. Outcomes are standardized to facilitate comparisons across regressions. For firm performance, we consider total employment, sales, and value added (both in levels and per-worker terms). Across all outcomes, point estimates are statistically insignificant and lie within 5% of a standard deviation.¹² This rules out improved firm performance as an explanation for the earnings gains we document.

Even if firm performance is unaffected, flexible arrangements might have raised employees' bargaining power, raising wages within the firm. To assess this, we examine changes in the earnings of mothers' coworkers—both all and male coworkers—and find no significant effects. This rules out changes in firm-level rent sharing as a driver of our results. The fact that mothers' earnings increase relative to their coworkers' confirms that we are capturing the effect of the contractual arrangement rather than a general increase in employees' bargaining power.

These results complement the placebo evidence, which shows that gains are concentrated around childbirth, with no effects in the years before or after, periods in which

suggests that it provides a good approximation of the longer-run benefits of flexible work arrangements.

¹²This aligns with Basso et al. (2025), who finds limited effects of flexible arrangements on firm performance except for a narrow set of firms. Our estimates are not directly comparable, however, as we capture effects at firms where mothers are employed rather than the average firm.

we would expect to observe impacts if the results were driven by improvements in firm performance or employees' bargaining power.

5.3 Heterogeneity by Individual, Family, and Job Attributes

If flexible work primarily relaxes time constraints around childbirth, its effects should be largest for groups facing the most binding constraints. Figure 4 summarizes the 2SLS heterogeneity estimates, while Appendix Tables A.2, A.3, and A.4 report the corresponding first-stage, reduced-form, and 2SLS coefficients in detail for the earnings change between the year after and two years before childbirth. In Appendix Figure A.4, 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.


We begin with heterogeneity by the mother's own characteristics. Splitting mothers by pre-childbirth earnings reveals sharp differences in how much they benefit from flexibility. Both in absolute terms and relative to pre-birth earnings, effects on the earnings change around childbirth are substantially larger for mothers in the bottom 75% of the earnings distribution than for those in the top 25%. The first stage is positive for both groups and markedly stronger for the top 25%, so the larger effects for lower earners are not driven by stronger take-up. Rather, the 2SLS estimates imply that a given increase in flexibility translates into much larger labor-market gains for lower-earning mothers. This is consistent with lower-earning mothers adjusting along more discrete labor supply margins—such as part-time work or temporary withdrawal from employment—that flexible arrangements help relax.

Age is informative about the role of labor market attachment and adjustment costs at the time of childbirth. The conceptual framework predicts stronger effects for younger mothers, whose labor market attachment is less established and for whom rigid schedules are particularly costly. The estimates are consistent with this. In the reduced form, absolute earnings gains are similar across age groups: mothers below 30 gain about 610 euro, while those aged 36 and above gain about 700 euro. However, the first stage is significantly smaller for the youngest group, implying a substantially larger 2SLS effect. Relative to pre-birth earnings, the difference between mothers below 30 and those aged 36 and above amounts to about 40 percentage points in the IV, statistically significant in both specifications. That younger mothers achieve similar level gains despite lower baseline earnings suggests larger adjustments along discrete labor supply margins such as part-time work and employment attachment.


Education points in the same direction. In levels, the 2SLS estimates are similar for mothers without college and for college-educated mothers, and the difference is not statistically significant. Yet because college-educated mothers have substantially higher baseline earnings, the relative effect is larger for non-college mothers. The percentage-

Figure 4: Heterogeneity in the effect of WFH on mothers' earnings around childbirth

(A) Effects in levels



(B) Effects as a share of pre-childbirth earnings



Notes. This figure reports 2SLS coefficients and their 95% confidence intervals from pre-post heterogeneous DID specifications for the earnings change around childbirth. Panel A reports estimates in euro. Panel B rescales the estimates by the subgroup-specific pre-2020 mean of pre-childbirth earnings, measured two years before the childbirth year. The heterogeneity dimensions are grouped into mother's characteristics, household characteristics, and job characteristics. Mother's characteristics include whether the mother's pre-childbirth earnings are in the top quartile of the distribution, the mother's age, and the mother's highest educational attainment; education is observed only for mothers matched to COB data. Household characteristics include whether the father's pre-childbirth earnings are in the top quartile of the distribution and whether the dominant earner in the household before childbirth is the mother or the father. Job characteristics include commuting status and exposure to greedy jobs. A commuter is a mother who works in a municipality different from the one where she lives. The greedy-jobs measure is constructed from EU-SILC data for Italy as the occupation-level elasticity of annual gross employee earnings with respect to hours worked, estimated on college-educated full-time employees aged 25–64; occupations are observed at the 2-digit level.

scaled IV difference amounts to about 11 percentage points in favor of mothers without college. This again suggests that flexibility is especially valuable for mothers whose labor supply is more exposed to discontinuous adjustments around childbirth.

We next turn to household characteristics. If household-level constraints are more binding in lower-income households, mothers should benefit more from flexibility when their partner is a lower earner. In such households, the opportunity cost of father-provided childcare is lower, making it easier to reallocate childcare toward fathers when mothers' constraints are relaxed. The estimates are consistent with this prediction. In levels, the 2SLS effects are similar across groups—about 3,400 euros for mothers whose partners are in the bottom 75% of the fathers' earnings distribution and about 3,700 euros for the top 25%. Relative to pre-birth earnings, however, the difference is statistically significant and amounts to about 8 percentage points in favor of mothers with lower-earning fathers. This is suggestive evidence that flexibility is more valuable in lower-income households, consistent with household-level constraints playing a central role in the motherhood penalty.

The dominant-earner split yields sharper evidence. Effects are much larger when the father is the main earner in the household before childbirth than when the mother is.¹³ The 2SLS estimate is about 4,700 euro when the father is the dominant earner, compared with roughly 2,500 euro when the mother is. The difference is large and precisely estimated both in levels and relative to pre-birth earnings. This pattern is consistent with the idea that childbirth generates larger labor supply distortions in households where fathers are the primary earners, because mothers are more likely to absorb the adjustment to new time constraints. Flexibility can attenuate these distortions by making it easier for mothers to remain attached to work and preserve hours after childbirth.

Finally, job characteristics provide direct evidence on the mechanisms through which flexibility matters. We first study commuting status, defined by whether a mother works in a different municipality from her residence. If the motherhood penalty reflects a binding time constraint, effects should be stronger for mothers who commuted before childbirth, since commuting absorbs time that flexible work can directly save. Both reduced-form and 2SLS effects are larger for commuters: the reduced-form effect is about 630 euro for commuters versus 300 euro for non-commuters, and the 2SLS effect about 4,600 versus 3,000 euro. The first stage is similar across groups, if anything slightly smaller for commuters, indicating that the heterogeneity reflects differences in the payoff from flexibility rather than in WFH take-up. The difference is statistically significant both in levels and in terms of percentage-scaled difference. Overall, these results are consistent with commuting as an additional driver of post-childbirth time constraints that raises the labor cost of childbirth.

We also find heterogeneity by exposure to greedy jobs. To measure greediness, we

¹³In our sample, 30% of mothers out-earn their partner in the private sector.

follow Goldin (2014) and construct an occupation-level index using pooled EU-SILC data for Italy and estimate, at the 2-digit occupation level, how steeply log annual earnings rise with log hours worked among college-educated full-time employees aged 25–64.¹⁴ We classify an occupation as greedy if the elasticity of earnings to hours is above 1. Mothers in greedier occupations display much larger gains from flexibility than mothers in less greedy jobs. The 2SLS estimate is about 3,200 euro in greedy jobs, compared with roughly 1,400 euro in non-greedy jobs, and the difference in levels is statistically significant. The relative difference is less precisely estimated, reflecting both smaller sample size and higher baseline earnings in greedier occupations, but the overall pattern is consistent: when occupations reward long and inflexible hours more strongly, access to flexible work has a larger payoff around childbirth. This is consistent with a mechanism in which flexibility attenuates the penalty generated by steep earnings-hours schedules.

Taken together, the heterogeneity results point to a common conclusion: flexible work generates the largest gains precisely where workplace rigidity is most costly. This holds across mother characteristics, household structure, and job characteristics. The results by father’s earnings and dominant-earner status support the view that the relevant constraint operates at the household level, while those by commuting and occupation point to job-specific sources of rigidity. These findings are consistent with a framework in which time and organizational rigidities are a key driver of the labor cost of motherhood.

6 Fathers and Household-level Constraints

6.1 Fathers’ Labor Market Outcomes

As a first step toward understanding household-level consequences, 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 his industry-location 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 5, for all fathers and for those in households where both partners are employed at childbirth.¹⁵

¹⁴Specifically, on the pooled EU-SILC Italy sample we estimate

$$\log Y_{iot} = \alpha + \beta F_i + \gamma_o + \delta_o F_i + (\eta + \kappa_o) \log H_{it} + \lambda \log M_{it} + \rho_t + p(\text{age}_{it}) + \varepsilon_{iot},$$

where Y_{iot} is annual gross employee income, F_i is a female dummy, H_{it} denotes usual weekly hours, M_{it} is a full-time-equivalent measure of months worked, ρ_t are year effects, and $p(\text{age}_{it})$ is a fourth-order polynomial in age. We estimate the regression on Italian employees aged 25–64 with college education, restricting to a full-time, full-year proxy sample. Occupation-specific greediness is then measured by the estimated occupation-specific elasticity of annual earnings with respect to hours, $\eta + \kappa_o$. We classify as greedy an occupation with $\eta + \kappa_o > 1$. Since EU-SILC reports occupations only at the 2-digit level, this is a coarser version of Goldin’s measure.

¹⁵The descriptive statistics for fathers in households are reported in column (4) of Table 1. The first

A well-established finding in the literature is that childbirth has limited consequences for fathers' earnings, a pattern we confirm in Italian data in Figure A.3. Consistent with this, exposure to remotable work has small and largely insignificant effects on fathers' earnings around childbirth, with only a mildly positive and marginally significant effect in 2021 for fathers in double-income households. Since fathers face little earnings penalty, flexible work has limited scope to generate gains.

The absence of earnings responses among fathers is informative about the mechanism: any impact of fathers' remotability on household outcomes operates not through increases in fathers' labor supply but through changes in intra-household time allocation.

6.2 Fathers' Remotability and Mothers' Motherhood Penalty

We next examine whether mothers' motherhood penalty depends on their partner's job remotability by augmenting our baseline specification to include both mothers' and fathers' exposure. This household-level specification jointly includes own and partner exposure, along with the corresponding levels for both individuals. As a result, the coefficient on fathers' exposure is identified conditional on mothers' own exposure, rather than driven by assortative matching. Table 1 reports descriptive statistics for the household-linked sample in columns (3) and (4).

Figure 6 shows that both exposures matter. As in our baseline analysis, a mother's own exposure increases her earnings and weeks worked change around childbirth (panels A and B), though the magnitude is slightly smaller once father's exposure is included, reflecting the positive correlation between the two exposures.¹⁶ We find an improvement of up to 700 euro (about 4% of pre-birth earnings). Father's exposure independently improves mothers' earnings by up to about 630 euro (roughly 3.6%)—a magnitude comparable to the direct effect of mothers' own remotability.¹⁷


As an additional placebo exercise, we estimate the same household-level specification for couples observed away from childbirth. To identify partners, we focus on women two years before childbirth and link them to their future partner, defining a placebo childbirth year as the actual year minus two. In this sample, neither maternal nor paternal exposure has detectable effects on mothers' earnings or weeks worked change around the placebo birth year (Figure A.6). This suggests that the main household-level results do not simply capture generic advantages of high-exposure couples, such as better jobs, more family-friendly employers, or assortative matching into flexible occupations. Instead, the effects emerge only around childbirth, when household time constraints become more binding.

stage is shown in Figure A.5.

¹⁶As Table A.1, panel D, shows, mothers matched to their partners also display slightly smaller effects of own remotability on earnings changes compared to the full sample of working mothers.


¹⁷Figure A.5 shows that the first stage for fathers is similar to that of mothers, so reduced-form estimates are approximately comparable one-to-one. Table A.5 summarizes the estimates in Panel A.

Figure 5: 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 6: Earnings and weeks of mothers, exposure to both own and father's remotability, 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) 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.

The absence of effects in the placebo sample therefore supports the interpretation that partner-side flexibility matters because it relaxes household constraints in the post-birth period.

Including in our specification the interaction between maternal and paternal exposure further addresses the concern that the partner effect reflects assortative matching among high-exposure couples. The interaction term is insignificant for earnings changes and negative and sizable for changes in weeks worked (Table A.5, panel B). This indicates that the effect of fathers' exposure is not concentrated among couples where both spouses are highly exposed. If anything, for weeks worked the effect is stronger when mothers have lower exposure, consistent with a household-constraint interpretation in which fathers' flexibility partly substitutes for mothers' limited flexibility.

These results are consistent with household-level constraints playing an important role in the labor cost of motherhood. A purely mother-side explanation—such as mother-specific career ladders, contracting frictions, or discrimination—would imply a more prominent role for mothers' own flexibility, with fathers' flexibility playing a more limited role. By contrast, the comparable magnitudes we estimate are easier to reconcile with a framework in which the motherhood penalty reflects a household time constraint interacting with rigid workplace arrangements for *both* spouses.¹⁸ When fathers' flexibility increases, they may be able to reallocate time toward childcare—possibly out of leisure, given the limited earnings responses—thereby easing mothers' time constraints and supporting their earnings. Overall, the evidence suggests that the motherhood penalty is not only a mother-specific labor market phenomenon, but also reflects joint household constraints interacting with workplace rigidity.

7 Flexible Arrangements and Fertility

We next examine whether flexible work arrangements affected fertility decisions among working women, testing whether increased availability of remote work led to higher child-birth rates. We focus first on the intensive and extensive margins of fertility, then we perform a back-of-the-envelope estimate of the impact of the diffusion of WFH on the total fertility rate. We conclude quantifying aggregate effects at the sector–location level.

Micro-level Evidence. Our conceptual framework links fertility to the motherhood 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 at least in the short-run. The effect should be strongest on the intensive margin—higher-order fertility—where the decision depends primarily on whether existing arrangements can absorb ad-

¹⁸This interpretation is also consistent with [Bang \(2022\)](#), who shows that spousal flexibility can be an important determinant of married women's labor supply around childbirth.

ditional 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. We therefore expect a stronger fertility response at higher parity than on the extensive margin, unless flexible arrangements are sufficient to offset the fixed costs of a first child.

We first focus on the intensive margin, constructing two samples from the UCA data that capture each woman’s complete fertility history, including births outside periods of employment. The *placebo* sample 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 (2016 or 2020), pool the two panels, and estimate:

$$F_{it} = \alpha_{s(i)} + \lambda_t + \sum_{k=2017}^{2024} \beta_k \cdot I(t = k) \cdot \text{Remotable}_{s(i)} + \sum_{k=2017}^{2024} \gamma_k \cdot I(t = k) \cdot X_{\bar{s}(i)} + \varepsilon_{ist}, \quad (4)$$


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

Figure 7 Panel A shows no evidence that fertility dynamics in the placebo sample are predicted by sector-level remotability. 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 about 0.011 in 2021 and 0.012 thereafter, corresponding to about 1.3%–1.5% of the average cumulative births in the placebo sample, and the effect remains approximately constant over the four years of the panel. Instrumenting WFH contract take-up with sector-location remotability, we find that holding such a contract raises cumulative fertility among incumbent mothers by about 0.12 children (Table 3). This is a sizable 15% increase relative to the placebo sample mean of 0.80, consistent with flexible work reducing the marginal cost of an additional birth.


We also examine the role of fathers’ remotability by including both partners’ exposure in the same regression. In contrast to what we find for mothers’ earnings, fathers’ remotability has no significant effect on household cumulative fertility (Figure A.7), suggesting that the fertility response operates primarily through mothers’ own work arrangements rather than through intra-household time reallocation.

Figure 7: Fertility outcomes and number of mothers, reduced-form estimates


(A) Incumbent mothers, cumulative number of children after the first



(B) First-time mothers, cumulative number of children



(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.

We obtain qualitatively similar results on the *extensive margin*, restricting attention to women without children by the end of the baseline year. Panel B shows no trend predicted by remotability in the placebo sample and a positive effect of remotability after the COVID shock: first births rise by 0.005 in 2021 (0.44% of the placebo average), with the effect declining over the following four years. The IV estimates in Table 3 point in the same direction: having a WFH contract raises cumulative fertility among women without prior children by about 0.04 children (3.0% of the placebo average). That flexible arrangements raise even first-birth rates suggests their gains are sufficient to offset part of the fixed costs of entering parenthood. The smaller magnitude relative to the intensive margin indicates, however, that these fixed costs are not fully absorbed, consistent with the remaining role of anticipatory career concerns and other non-economic factors relevant in the decision to become a mother.¹⁹ The declining pattern of the estimates further indicates that part of the observed increase reflects a timing effect, with childbirth occurring earlier rather than a rise in total fertility.

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

For a back-of-the-envelope estimate of the impact of remotability on the total fertility rate, we combine the IV estimates in Table 3 with age-specific changes in WFH exposure and population counts. Specifically, $\hat{\beta}_g$ is the corresponding IV coefficient from Table 3, where g indexes whether women are first-time mothers (without children) or incumbent mothers (with at least one child), ΔWFH_{ga} is the change in WFH exposure by age and group measured in Labor Force Survey microdata²⁰ between 2019 and the average over 2020–2024, P_{ga}^E is the number of baseline-employed women in each cell measured in administrative data, and N_a is the total female population by age from the National Statistical Institute (Istat). Formally, for each age a we compute the implied increase in births over the four-year horizon as $\Delta B_a^{(4)} = \sum_g \hat{\beta}_g \Delta WFH_{ga} P_{ga}^E$. The implied effect on the total fertility rate is then $\Delta TFR = \sum_a (\Delta B_a^{(4)} / 4) / N_a$. This exercise implies about 4,700 additional births per year under the observed rise in WFH, corresponding to an increase in the Italian TFR of 0.014, or 1.2% of its 2024 level. Under a counterfactual in which WFH rises to its estimated potential among workers whose jobs are compatible with remote work, the implied effect rises to about 15,200 births per year and 0.046 TFR points, equal to 3.9% of the 2024 TFR. Relative to the decline in Italian fertility since its peak in 2008—the highest level since the late 1980s—these magnitudes account for

¹⁹We cannot replicate the household-level analysis for first births since in our data we are not able to match couples in case they have no children.

²⁰We compute WFH exposure separately by age and group. Actual WFH exposure is the share of female employees working from home. Potential WFH exposure is obtained by assigning each worker an occupation-level measure of remote-work feasibility and averaging it within age-group cells. For the observed scenario, ΔWFH_{ga} is the change in actual WFH exposure between 2019 and the average over 2020–2024. For the potential-WFH scenario, it is the difference between potential WFH exposure and actual WFH exposure in 2019.

about 5.4 and 17.7 percent, respectively.²¹

Sector-location Total Fertility. We next quantify the aggregate fertility response by adapting specification (1) to the sector–location level, computing for each cell the number of mothers and constructing a balanced panel that includes zeros for unobserved combinations.

Panel C of Figure 7 reports results also distinguishing by age group: women below 26, aged 27–35, and above 35. Moving from the 10th to the 90th percentile of Remotable_s increases the total number of mothers in a sector–location by 0.44 (1.2% of the pre-2020 average), driven by the two older groups: mothers aged 27–35 increase by 0.29 and those aged 35+ by 0.12, relative to pre-COVID averages of 23.4 and 11.4, respectively. Table 3 reports IV estimates where the sector–location WFH take-up rate among working women aged 20–45 is instrumented with remotability: moving from 0 to 100% take-up of WFH contracts raises the number of mothers by about 2.4 in a sector–location (6.5% of the pre-2020 average). These results suggest that some women in older age groups chose to have children following the expansion of flexible work.

Whether this fertility increase represents a timing adjustment or a structural rise in total fertility cannot be determined with data covering only a few years after the COVID shock. However, even if it reflects anticipation of births, 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 Work and the Gender Earnings Gap

Having established that flexibility is a central driver of the motherhood penalty, we conclude by quantifying how the diffusion of flexible work arrangements affects the gender earnings gap. This gap widens over a cohort’s lifetime (Arellano-Bover et al., 2024; Goldin, 2014), with the motherhood penalty being one of its most important drivers (Goldin et al., 2024).

We assess the extent to which WFH opportunities may reduce this widening through a back-of-the-envelope exercise. We focus on cohorts born between 1974 and 1979 who reached the end of their fertile years before COVID-19. We select those who worked at least 156 weeks (*i.e.*, 3 years) before turning 30 and construct a balanced panel, assigning zero earnings to years in which individuals are not observed. As expected, the earnings

²¹The aggregate effect is split fairly evenly between the two margins. Under the observed change in WFH, the implied TFR increase is 0.0066 for incumbent mothers and 0.0075 for first-time mothers, equal to 1.56% of the 2024 second-order TFR and 1.29% of the 2024 first-order TFR, respectively. Under the potential-WFH scenario, the corresponding effects are 0.0256 and 0.0204, equal to 6.10% of the 2024 second-order TFR and 3.51% of the 2024 first-order TFR.

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


	First-stage	Reduced-form	2SLS	Mean outcome
<i>Panel A: Women with kids</i>				
Post 2020 cohort	0.098*** (0.005)	0.012*** (0.002)	0.121*** (0.019)	0.803
KP F-stat	382.013			
Observations	6,705,168	6,705,168	6,705,168	
<i>Panel B: Women without kids</i>				
Post 2020 cohort	0.112*** (0.006)	0.004*** (0.001)	0.036*** (0.011)	1.174
KP F-stat	297.090			
Observations	17,182,532	17,182,532	17,182,532	
<i>Panel C: Sector-location level fertility</i>				
Total	0.089*** (0.002)	0.216*** (0.048)	2.421*** (0.548)	37.366
Age 18-26	0.048*** (0.001)	0.011** (0.005)	0.228** (0.099)	2.558
Age 27-35	0.086*** (0.002)	0.155*** (0.030)	1.812*** (0.355)	23.422
Age 36+	0.095*** (0.002)	0.050** (0.020)	0.527** (0.216)	11.386
KP F-stat	2,586.275			
Observations	2,928,000	2,928,000	2,928,000	

Notes. Each column reports estimates from pre-post difference-in-differences specifications for fertility outcomes. The first-stage column reports the effect of remotability on the probability of holding a WFH contract in Panels A and B, and on the share of women with WFH contracts in the sector-location in Panel C. The reduced-form column reports the effect of remotability on the outcome, and the 2SLS column reports the implied effect of WFH. In Panels A and B, the outcome is the cumulative number of children over the four years of the post-COVID panel of women. Panel A refers to women with kids, Panel B to women without kids, and Panel C to sector-location level fertility, measured by the number of mothers overall and by age group. The last column reports the mean of the corresponding outcome: in Panels A and B this is the mean in the placebo sample, while in Panel C it is the pre-2020 mean. Standard errors are reported in parentheses, and KP F-stat denotes the Kleibergen-Paap weak-instrument statistic.

gender gap widens over the lifecycle. Figure 8 plots the female-to-male average earnings ratio from age 25 to 40. Starting at 89.4% at age 25, it declines steadily to 70.4% by age 40—a widening of 19.0 percentage points over the first 15 years of the career.

We construct three counterfactuals for the female-to-male earnings ratio: (i) removing the full motherhood penalty; (ii) incorporating the reduction in the motherhood penalty under current WFH adoption; and (iii) projecting the gap reduction if all remotable jobs

Figure 8: Gender gap in labor earnings over the life cycle under different motherhood 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 motherhood penalty” assumes that the motherhood penalty in annual labor earnings is null.

adopted WFH contracts.²² We calibrate the penalty reduction attributable to WFH using our estimates in Section 5.1.²³ The first counterfactual delivers the largest reduction, the second the smallest, and the third illustrates the potential for further gains from a broader WFH diffusion.

The results highlight two main insights. First, consistently with Landais et al. (2025), the motherhood 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. Moreover, because a share of women already have children by age 25, removing the motherhood penalty would also shift the earnings ratio upward by about 2.8 percentage points.

Second, WFH arrangements attenuate the widening of the gender earnings gap. Under

²²Both actual and potential WFH adoption are measured using the Italian Labour Force Survey. Actual WFH adoption is computed as the share of female employees working from home in 2020, by age. Potential WFH adoption is computed by assigning to each worker the occupation-level share of jobs that can be performed remotely, based on the Dingel and Neiman (2020) occupational classification, and then averaging this measure across female employees in 2020, by age.

²³Specifically, we use the estimate of ρ from the mothers-non-mothers comparison calibration.

current adoption rates, the female-to-male ratio improves by 2.2 percentage points relative to the observed path, equivalent to 10.7% of the lifetime widening. Full adoption across all remotable jobs would improve the female-to-male ratio at age 40 by 6.1 percentage points and reduce the lifecycle widening by 5.6 percentage points, corresponding to 29.4% 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. This depends on the limited number of jobs that can be performed remotely.

Appendix Figure A.8 shows that, for our reference cohorts, the correlation between sector remotability and the share of mothers was positive but negligible, indicating that mothers were not disproportionately concentrated in remotable sectors prior to COVID-19. By contrast, as documented in Section 7, post-COVID fertility has increased precisely in more remotable sectors. Younger cohorts of women therefore tend to have children in sectors more exposed to WFH, implying that the lifecycle earnings gap reduction quantified in Figure 8 is likely a lower bound for the gains that younger generations will experience as flexible work continues to spread.

9 Conclusions

This paper provides causal evidence that rigid workplace arrangements are a key driver of the motherhood 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 for the jobs that can be performed remotely.

Our results also shed light on the mechanism. Fathers’ earnings do not respond to greater exposure to remote work around childbirth, yet fathers’ remotability reduces mothers’ earnings losses by a magnitude comparable to mothers’ own exposure. This pattern is difficult to reconcile with explanations based solely on mother-specific dynamics, discrimination, or fixed gender norms, and instead points to a binding household-level time constraint interacting with rigid job structures. Consistent with this interpretation, greater availability of flexible contracts lowers the labor-market cost of children and increases fertility, highlighting how workplace organization shapes both earnings trajectories and family decisions.

Expanding access to flexible work arrangements may be more effective at reducing the motherhood penalty than paternal leave policies that entail earnings penalties for fathers (Andresen and Nix, 2025), as they promote paternal involvement while lowering the labor-market cost of children for mothers. In Appendix B, we show that, accounting for earnings gains and reduced parental leave subsidies, the social value of a WFH contract is about 6,428 euros. The optimal subsidy to WFH adoption equals the contract’s social

value multiplied by $\varepsilon/(1+\varepsilon)$, where ε is the take-up elasticity, highlighting the importance of estimating the responsiveness of WFH adoption to financial incentives.

Finally, the post-COVID diffusion of WFH has the potential to slow the life-cycle widening of the gender earnings gap. While WFH alone cannot eliminate gender inequality because many jobs cannot be performed remotely, our evidence indicates that contractual rigidity is a first-order component of the motherhood penalty. 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. Working Paper 21250, CEPR.
- Albanesi, S. (2026). Inflexibility and the Stalled Progress of American Women in the Labor Market. Working Paper 21426, CEPR.
- 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). Parenthood 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*.
- Bang, M. (2022). Job Flexibility and Household Labor Supply: Understanding Gender Gaps and the Child Wage Penalty. Working paper, University of Cambridge.
- 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.
- Coibion, O., Gorodnichenko, Y., and Weber, M. (2025). The cost of the covid-19 crisis: Lock-downs, macroeconomic expectations, and consumer spending. *Journal of Economic Behavior & Organization*, 229:106846.
- 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.
- Guerrieri, V., Lorenzoni, G., Straub, L., and Werning, I. (2022). Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? *American Economic Review*, 112(5):1437–74.
- 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 Søgaaard, 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)).
- Landais, C., Lundborg, P., Pazen, T., and Vikstrom, J. (2025). Gender Without Children. Working Paper.
- 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.
- Lu, S., Zang, E., and Musick, K. a. (2025). A Remote Work Baby Boom? Remote Work and Fertility in the United States. Working Paper.
- Mas, A. and Pallais, A. (2017). Valuing Alternative Work Arrangements. *American Economic Review*, 107(12):3722–3759.

- Monte, F., Porcher, C., and Rossi-Hansberg, E. (2026). Remote Work and City Structure. *American Economic Review*. forthcoming.
- 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.

Online Appendix to

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

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

Additional Tables and Figures

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

	(1) Earnings	(2) Net earnings	(3) Net earnings t+2	(4) Weeks	(5) Net weeks	(6) PL weeks	(7) Employed	(8) Employed full-year	(9) Part-time	(10) Promoted post
<i>Panel A: Baseline</i>										
Remotable × Post	978.57*** (149.84)	512.86*** (105.81)	501.45*** (110.20)	0.53*** (0.14)	0.70*** (0.16)	-0.71*** (0.10)	-0.001 (0.002)	0.015*** (0.004)	-0.022*** (0.004)	0.003*** (0.001)
Dep. var. mean	11395	-5105	-3257	25.6	-9.39	11.6	.86	.329	.457	.00192
Observations	770,625	770,625	679,226	770,625	770,625	770,625	770,625	770,625	661,562	661,562
<i>Panel B: Location-by-year fixed effects</i>										
Remotable × Post	946.90*** (141.14)	473.55*** (111.77)	464.71*** (123.39)	0.62*** (0.15)	0.52*** (0.17)	-0.58*** (0.10)	0.002 (0.003)	0.016*** (0.004)	-0.017*** (0.004)	0.002*** (0.001)
Dep. var. mean	11395	-5105	-3257	25.6	-9.39	11.6	.86	.329	.457	.00192
Observations	770,625	770,625	679,226	770,625	770,625	770,625	770,625	770,625	661,562	661,562
<i>Panel C: Wild bootstrap cluster at year level</i>										
Remotable × Post	978.57*** (192.71)	512.86*** (161.65)	501.45*** (109.59)	0.53*** (0.18)	0.70** (0.32)	-0.71*** (0.07)	-0.001 (0.003)	0.015*** (0.004)	-0.022*** (0.002)	0.003*** (0.001)
Dep. var. mean	11395	-5105	-3257	25.6	-9.39	11.6	.86	.329	.457	.00192
Observations	770,625	770,625	679,226	770,625	770,625	770,625	770,625	770,625	661,562	661,562
<i>Panel D: Mothers in household</i>										
Remotable × Post	1123.04*** (184.68)	452.00*** (127.69)	520.15*** (140.32)	0.60*** (0.17)	0.49** (0.19)	-0.78*** (0.12)	-0.000 (0.003)	0.019*** (0.005)	-0.026*** (0.005)	0.004*** (0.001)
Dep. var. mean	11752	-5655	-3551	26.5	-10.3	12.1	.867	.349	.429	.00197
Observations	405,844	405,844	350,834	405,844	405,844	405,844	405,844	405,844	350,485	350,485

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 adds province-by-year fixed effects. Panel C reports baseline estimates with standard errors clustered also at the year level. Due to the small number of year clusters, in this case we compute standard errors with a wild bootstrap procedure. Panel D reports estimates for the subsample of mothers living in households in which both parents are employed at the time of childbirth. The outcomes, measured with a one-year lead unless otherwise specified, are: annual earnings; annual earnings net of pre-childbirth earnings; total full-time equivalent weeks net of pre-childbirth earnings; total full-time equivalent weeks worked; total full-time equivalent weeks net of pre-childbirth weeks; parental leave duration in weeks; an indicator for being employed; an indicator for being promoted (conditional on employment) to a managerial position, from either blue- or white-collar jobs. Standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: Heterogeneity by mother's characteristics: Reduced form and 2SLS static difference-in-differences

	First-stage	Reduced-form	2SLS	Pre-2020 earnings at $t - 2$
<i>Panel A: Mother's earnings</i>				
Bottom 75%	0.07*** (0.01)	784.25*** (118.33)	10132.76*** (1578.06)	11179.29
Top 25%	0.18*** (0.01)	431.73*** (154.61)	4034.14*** (953.62)	32248.02
KP F-stat	106.364			
Difference: Bottom 75% vs Top 25%		352.519** [2.226]	6098.614*** [5.563]	
Difference (%): Bottom 75% vs Top 25%		0.057*** [5.692]	0.781*** [6.443]	
Observations	770,625	770,625	770,625	
<i>Panel B: Age</i>				
Under-30	0.06*** (0.01)	612.84*** (133.93)	6646.41*** (1250.93)	10755.40
30-35	0.14*** (0.01)	379.40*** (115.26)	3456.25*** (801.49)	17036.26
36+	0.14*** (0.01)	702.56*** (187.91)	4697.78*** (1093.72)	21197.43
KP F-stat	100.022			
Difference: Under-30 vs 36+		-89.720 [-0.446]	1948.631* [1.790]	
Difference (%): Under-30 vs 36+		0.024* [1.788]	0.396*** [4.115]	
Observations	770,625	770,625	770,625	
<i>Panel C: Education</i>				
No college	0.06*** (0.01)	264.06*** (94.31)	3978.06*** (1167.26)	13730.10
College	0.16*** (0.01)	553.73*** (152.71)	3616.19*** (923.29)	19812.72
KP F-stat	115.927			
Difference: College vs no college		289.669* [1.789]	-361.875 [-0.399]	
Difference (%): College vs no college		0.009 [0.948]	-0.107 [-1.639]	
Observations	716,227	716,227	716,227	

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. Panel A reports heterogeneity by whether the mother's pre-childbirth earnings are in the top quartile of the distribution. Panel B reports heterogeneity by the mother's age. Panel C reports heterogeneity by the mother's highest educational attainment. Education is observed only for mothers matched to COB data.

Table A.3: Heterogeneity by household characteristics: Reduced form and 2SLS static difference-in-differences

	First-stage	Reduced-form	2SLS	Pre-2020 earnings at $t - 2$
<i>Panel A: Father's earnings</i>				
Bottom 75%	0.11*** (0.01)	353.35*** (123.69)	3374.14*** (1100.28)	14971.46
Top 25%	0.16*** (0.01)	607.34*** (179.04)	3688.38*** (1041.02)	24619.38
KP F-stat	148.331			
Difference: Bottom 75% vs Top 25%		-253.989 [-1.605]	-314.240 [-0.459]	
Difference (%): Bottom 75% vs Top 25%		-0.001 [-0.137]	0.076 [1.598]	
Observations	405,843	405,843	405,843	
<i>Panel B: Dominant earnings</i>				
Mother higher earnings	0.15*** (0.01)	223.59 (153.30)	2541.20*** (961.65)	22725.52
Father higher earnings	0.12*** (0.01)	684.17*** (141.56)	4744.99*** (1042.43)	14667.29
KP F-stat	166.353			
Difference: Mother higher vs Father higher		-460.579*** [-3.151]	-2203.792*** [-3.556]	
Difference (%): Mother higher vs Father higher		-0.037*** [-4.334]	-0.212*** [-4.782]	
Observations	405,843	405,843	405,843	

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. Panel A reports heterogeneity by whether the father's pre-childbirth earnings are in the top quartile of the distribution. Panel B reports heterogeneity by whether the dominant earner in the household before childbirth is the mother or the father.

Table A.4: Heterogeneity by job characteristics: Reduced form and 2SLS static difference-in-differences

	First-stage	Reduced-form	2SLS	Pre-2020 earnings at $t - 2$
<i>Panel A: Commuting</i>				
Non-commuter	0.14*** (0.01)	301.43* (159.74)	3012.94*** (950.64)	15366.07
Commuter	0.13*** (0.01)	631.53*** (99.44)	4564.10*** (786.82)	17217.59
KP F-stat	179.071			
Difference: Commuter vs non-commuter		330.098** [2.277]	1551.160*** [2.635]	
Difference (%): Commuter vs non-commuter		0.017* [1.829]	0.069* [1.800]	
Observations	770,625	770,625	770,625	
<i>Panel B: Greedy jobs</i>				
Non-greedy job	0.11*** (0.01)	93.63 (115.66)	1434.49 (950.90)	13676.84
Greedy job	0.14*** (0.01)	498.41*** (130.40)	3245.91*** (879.26)	19574.39
KP F-stat	157.693			
Difference: Greedy vs non-greedy job		404.783*** [2.873]	1811.420** [2.464]	
Difference (%): Greedy vs non-greedy job		0.019** [2.128]	0.061 [1.194]	
Observations	591,209	591,209	591,209	


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. Panel A reports heterogeneity by commuting status. A commuter is a mother who works in a municipality different from the one where she lives. Panel B reports heterogeneity by exposure to greedy jobs. The greedy-jobs measure is constructed from EU-SILC data for Italy as the occupation-level elasticity of annual gross employee earnings with respect to hours worked, estimated on college-educated full-time employees aged 25–64; occupations are observed at the 2-digit level.

Table A.5: Household-level results

	Net earnings (1)	Net weeks (2)
<i>Panel A: Mother and father exposure</i>		
Remotable ^m × Post	318.36** (124.87)	0.28 (0.21)
Remotable ^f × Post	283.41*** (101.96)	0.44*** (0.15)
Mean Dep. Var	-5655	-10.3
Observations	405,843	405,843
<i>Panel B: Mother, father, and interaction exposure</i>		
Remotable ^m × Post	272.31** (135.06)	0.51** (0.24)
Remotable ^f × Post	236.74* (122.43)	0.73*** (0.22)
Remotable ^m × Remotable ^f × Post	77.56 (163.27)	-0.41** (0.21)
Mean Dep. Var	-5655	-10.3
Observations	405,843	405,843


Notes. This table reports reduced-form coefficients from static difference-in-differences specifications in which mothers' and fathers' exposure to remotability are interacted with a post-2020 indicator. Panel A includes only the mother's and father's sector-location exposure, while Panel B adds their interaction. The dependent variables are the change in mothers' earnings around childbirth, net of pre-childbirth earnings, and the corresponding change in full-time equivalent weeks worked. Remot^m denotes the mother's sector-location exposure, while Remot^f denotes the father's sector-location exposure, both measured at childbirth. Post is a dummy for years after 2019. Standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A.1: Log change of labor market outcomes, dynamic reduced form estimates




Notes. This figure reports reduced-form coefficients and 95% confidence intervals from estimating equation (1) for changes in log annual earnings, weekly earnings, and weeks worked between two years before childbirth and the year after childbirth. Confidence intervals are based on standard errors clustered at the 4-digit sector-by-local-labor-market level.

Figure A.2: Placebo analysis of earnings and weeks changes




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, for placebo samples of women without children who are assigned a placebo childbirth year. 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 the assigned placebo childbirth year. The two series compare women assigned a placebo childbirth year two years before an actual future childbirth with women assigned a placebo childbirth year three years after an actual past childbirth. Because both placebo samples are constructed from women in the main analysis window, the first placebo can only be observed up to placebo year 2021, since later years would require post-placebo outcomes beyond 2024. By contrast, the second placebo starts in 2019, three years after the first childbirth cohort in 2016, and runs through 2023.

Figure A.3: The motherhood penalty




Notes. The graphs display event-time coefficients estimated from a regression of annual earnings on year-to-childbirth dummies, calendar year dummies and individual fixed effects ($Y_{ist}^g = \theta_i + \sum_{j \neq -1} \alpha_j^g \cdot \mathbf{1}[j = t] + \sum_y \gamma_y^g \cdot \mathbf{1}[y = s] + \nu_{ist}^g$), as a percentage of the counterfactual outcome in the absence of children (i.e., $P_t^g \equiv \hat{\alpha}_t^g / E[Y_{ist}^g | t]$). Panel A reports estimates for mothers and fathers; Panel B reports estimates for mothers and placebo non-mothers separately; the specification is the same. Placebo non-mothers are identified as in [Casarico and Lattanzio \(2023\)](#). The figure also reports a “child penalty,” defined as the percentage by which mothers fall behind the relevant comparison group after having children: $P_t \equiv (\hat{\alpha}_t^m - \hat{\alpha}_t^c) / E[Y_{ist}^c | t]$, where c denotes fathers in Panel A and placebo non-mothers in Panel B. The short- and long-run child penalties correspond to event time 1 and 10, respectively. In both panels, the coefficients are estimated unconditional on employment status, imputing zeros when the worker is not employed, and shaded areas represent 95% confidence intervals based on standard errors clustered at the individual level. Panel A is computed on the universe of INPS data, extended from 1997, and is estimated on a balanced sample of workers who had their first child between 2003 and 2024 with age at childbirth between 20 and 45; the information on childbirth is taken from the Universal Childcare Allowance database. Panel B is computed on a sample of INPS data for workers born on 24 dates of the year and is 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 information on childbirth is taken from maternity leave events.

Figure A.4: 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, and fathers' earnings quartile. Confidence intervals are based on standard errors clustered at the 4-digit sector and local labor market level.

Figure A.5: 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.6: Earnings and weeks of placebo mothers, exposure to both own and father's remotability, 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) 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. The reference population is composed of placebo mothers, defined as women who do not have a child at time t but will give birth two years later.

Figure A.7: Cumulative number of children after the first one for incumbent mothers, exposure to both own and father's remotability, 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) from the estimation of equation (4), augmented with the remotability of the father's job. The regressions are estimated on the 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. A balanced panel is constructed to track fertility outcomes over the five years following the baseline. The outcome is the cumulative number of children.

Figure A.8: Correlation between the share of mothers at 40 years old and remotability at the sector level



Notes. This figure reports a binned scatter plot of the relationship between 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 and remotability at the sector level. The 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:

$$\ell_i = 1 - T(k) - h_i - \phi_i(k, a),$$

where $\phi_i(k, a)$ captures coordination burdens from combining work and childcare. Flexible arrangements lower these burdens, so $\phi(\cdot, a') < \phi(\cdot, a)$ for $a' > a$. If one wants to add commuting or other job-level fixed time costs that do not depend on children or hours, they can be introduced as an additive term $\chi(a)$ in leisure or absorbed into $T(k)$ without changing the main child-penalty logic. Flexibility is especially valuable at higher parity: $\partial^2 \phi(k, a) / \partial k \partial a < 0$.²⁴

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

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

Assume $1 - T(k) - \phi(k, a) > 0$, so the optimum is interior and satisfies $h^*(k, a) > 0$ and $\ell^*(k, a) > 0$. The optimal hours and earnings are

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

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

Proposition 1 (Flexibility reduces the child penalty). *Greater remotability increases $E(k, a)$. It decreases $\Delta(k, a)$ if flexibility lowers child-related coordination costs sufficiently more at higher parity, so that $\phi(k, a) - \phi(0, a)$ falls with a .*

Proof. From the closed form,

$$\Delta(k, a) = \frac{w(T(k) - T(0) + \phi(k, a) - \phi(0, a))}{1 + \alpha}.$$

Hence $E(k, a)$ rises when $\phi(k, a)$ falls. The child penalty falls when flexibility reduces the child-specific component $\phi(k, a) - \phi(0, a)$. $\square$

A.3 Heterogeneous effects across mothers

A.3.1 High- vs low-earnings mothers

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

²⁴This assumption is used only in the fertility result (Section A.5).

earnings is

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

Since lower-earning mothers have larger coordination burdens, the denominator is smaller and the proportional gain is mechanically larger for them when flexibility relaxes these frictions. 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), and let the current wage be a common w . Current post-birth earnings are

$$E_0(k, s, a) = w \frac{1 - T(k) - \phi(k, s, a)}{1 + \alpha}.$$

Career accumulation affects wage growth between the immediate post-birth period and a later period:

$$g(s, a) = G(H(s) - \delta(a)), \quad G' > 0, G'' < 0,$$

where $H(s)$ is accumulated human capital with $H'(s) > 0$ and $\delta(a)$ captures career disruption with $\delta_a < 0$. Let next-period wage be

$$w_1(s, a) = w g(s, a).$$

To keep the extension parsimonious, assume next-period hours return to a common benchmark $\bar{h} > 0$, so

$$E_1(s, a) = w_1(s, a) \bar{h}.$$

Total post-birth value is $V(k, s, a) = E_0(k, s, a) + \beta E_1(s, a)$. Allowing current wages to rise with seniority would not change the qualitative mechanism, but is omitted here for simplicity.

Proposition 2 (Age heterogeneity). *Suppose (i) flexibility reduces current frictions more strongly at earlier career stages: $\phi_{as} > 0$; (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

$$\frac{\partial E_0}{\partial a} = \frac{w}{1 + \alpha} (-\phi_a(k, s, a)) > 0.$$

By $\phi_{as} > 0$, differentiating with respect to s gives $\frac{\partial}{\partial s} \left(\frac{\partial E_0}{\partial a} \right) < 0$. The future earnings effect satisfies

$$\beta \frac{\partial E_1}{\partial a} = -\beta \bar{h} w G'(H(s) - \delta(a)) \delta_a > 0, \quad \frac{\partial}{\partial s} \left(\beta \frac{\partial E_1}{\partial a} \right) = -\beta \bar{h} w \delta_a G''(H(s) - \delta(a)) H'(s) < 0,$$

using $G'' < 0$, $H'(s) > 0$, and $\delta_a < 0$. Summing both terms gives the result. $\square$

Commuting distance. Households with longer commutes should benefit more from flexibility if on-site work carries additional fixed time costs. In the streamlined formulation, such costs can be absorbed into $T(k)$ or added as an additive term $\chi(a)$ in leisure. The comparative static is immediate: if longer commutes imply a larger baseline time requirement, then any arrangement that reduces this fixed burden relaxes the mother's time constraint more strongly and generates larger gains in post-birth hours and earnings.

Education. Education can proxy for several distinct margins. More-educated mothers may face lower baseline coordination costs because they hold more autonomous jobs, but they may also have steeper career ladders, so flexibility has greater value in preventing dynamic losses. In the notation above, education can shift either the current-period friction $\phi(k, a)$ or the dynamic term $\delta(a)$. The direction of heterogeneity is therefore not mechanical and depends on which margin dominates empirically.

Greedy jobs. The linear earnings specification can be extended to sectors where earnings are convex in hours. Let period-0 earnings be $y(h)$ with $y'(h) > 0$ and $y''(h) > 0$, so that long hours are disproportionately rewarded. The mother's problem becomes

$$\max_{h \geq 0} \ln y(h) + \alpha \ln(1 - T(k) - \phi(k, a) - h).$$

The first-order condition is

$$\frac{y'(h)}{y(h)} = \frac{\alpha}{1 - T(k) - \phi(k, a) - h}.$$

Because $y''(h) > 0$, a given flexibility-induced relaxation of the time constraint has a larger effect on earnings when the mother is near the steep part of the earnings schedule. This provides a simple rationale for larger flexibility effects in “greedy” occupations: preserving hours after childbirth avoids a disproportionately large earnings loss.

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

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

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

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

subject to the childcare feasibility constraint $t_m + t_f = T(k)$.²⁵ The first-order conditions for market hours are

$$\alpha_i \frac{c}{\ell_i} = w_i, \quad i \in \{m, f\}, \quad (\text{A.1})$$

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

$$\frac{\alpha_m}{\ell_m} = \frac{\alpha_f}{\ell_f}, \quad (\text{A.2})$$

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 $w_f > w_m$, a marginal reallocation of childcare from mother to father reduces household consumption by w_f rather than 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 earnings and dominant-earner households. The household formulation also clarifies heterogeneity by fathers' earnings and by which spouse is the dominant earner. These are related but distinct margins. If $w_f > w_m$, comparative advantage pushes childcare and

²⁵The 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.

labor-supply adjustment toward mothers. Rigid arrangements then generate larger losses for them, so mother-side flexibility can have stronger effects in father-dominated households. A separate margin concerns the level of fathers' earnings: if fathers earn less in absolute terms, the opportunity cost of father-provided childcare is lower, so the household can more easily reallocate childcare toward them when mothers' constraints are relaxed. Thus stronger effects in father-dominated households and in low-father-earnings households need not contradict each other: the first works through relative earnings within the couple, while the second works through the absolute cost of shifting childcare to fathers.

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

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

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

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

so $\frac{\partial h_m^*}{\partial \bar{t}_f} = \frac{1}{1 + \alpha_f + \alpha_m} > 0$. Since $\bar{t}_f$ increases when ϕ_f falls, the result follows. $\square$

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 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

$$W(b; k) = V(k) + \beta \left(V(k+b) + B(k+b) - bF(k+b) \right),$$

where $B(\cdot)$ is the utility benefit from children and $F(k) > 0$ is the fixed cost of an additional birth, which decreases at higher parity ($F'(\cdot) < 0$). It chooses $b = 1$ whenever

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

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 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 $\partial^2 \phi(k, a) / \partial k \partial a < 0$, which makes flexibility reduce coordination costs more at higher parity. Hence $\Delta W(k)$ increases and the household is more likely to choose $b = 1$. $\square$

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 Details on WTP and MVPF computations

This Appendix provides a simple framework clarifying the assumptions under which our reduced-form estimates can be mapped into welfare formulas. It then provides details on the parameters calibration and the implementation of the formulas.

B.1 Micro-foundations for WTP and MVPF

Household. Let $z \in \{0, 1\}$ denote access to a WFH contract. A household solves

$$V(z) = \max_{h_m, h_f, t_m, t_f} u(c, \ell_m, \ell_f) \quad (\text{A.4})$$

subject to

$$c = w_m h_m + w_f h_f + R, \quad (\text{A.5})$$

$$\ell_i = 1 - (1 + \tau_i(z))h_i - t_i - \phi_i(z), \quad i \in \{m, f\}, \quad (\text{A.6})$$

$$t_m + t_f \geq T. \quad (\text{A.7})$$

Here h_i denotes market work, t_i childcare time, ℓ_i leisure, $\tau_i(z)$ a time wedge capturing commuting and schedule rigidity, and $\phi_i(z)$ a fixed coordination cost of combining work and care. WFH relaxes constraints if

$$\tau_i(1) \leq \tau_i(0), \quad \phi_i(1) \leq \phi_i(0). \quad (\text{A.8})$$

Let EV_H denote the equivalent variation associated with access to WFH. The relevant beneficiary-side welfare object is EV_H , not the observed earnings change, since labor supply is chosen endogenously.

Firm. A firm chooses whether to adopt WFH, $a \in \{0, 1\}$, and earns

$$\Pi(a; s) = q(a) - w(a) - \kappa(a) + sa, \quad (\text{A.9})$$

where $q(a)$ is output, $w(a)$ labor costs, $\kappa(a)$ the implementation cost of WFH, and s a subsidy paid conditional on adoption. Adoption occurs if $\Pi(1; s) \geq \Pi(0; s)$.

Assumptions and formula derivations: Suppose that: (i) WFH weakly expands the household feasible set by lowering time and coordination wedges; (ii) non-pecuniary effects of WFH are weakly positive; and (iii) firm surplus at the margin is negligible, $\Delta\Pi \approx 0$. The latter can be consequence of the fact that a change in adoption creates no first-order effects on firm surplus beyond the mechanical transfer receives. It is consistent with the absence of detectable effects on firm outcomes and coworkers' earnings in the data. Then household equivalent variation is conservatively approximated by after-tax earnings gains,

$$EV_H \approx (1 - \tau)\Delta Y, \quad (\text{A.10})$$

which implies that after-tax earnings gains are a money-metric proxy for household welfare. The fiscal externality arises from changes in earnings and parental leave:

$$FE = \tau\Delta Y + b_{PL}(-\Delta PL), \quad (\text{A.11})$$

where τ is the effective tax wedge on earnings and b_{PL} is the fiscal cost per week of parental leave.

The total welfare gains from an induced WFH contract are equal to the sum of the fiscal externality and the increase in household welfare. These gains correspond to the maximum

subsidy the government is willing to implement per induced contract. The latter therefore is:

$$s^{\max} = EV_H + FE \approx \Delta Y + b_{PL}(-\Delta PL). \quad (\text{A.12})$$

MVPF and the role of adoption elasticity. Let $p(s)$ denote the WFH adoption rate under subsidy s , and define the elasticity of adoption with respect to the subsidy as

$$\varepsilon = \frac{s}{p} \frac{dp}{ds}. \quad (\text{A.13})$$

For a small increase in the subsidy ds , total subsidy spending changes by

$$d(sp(s)) = p ds + s dp.$$

The first term is the mechanical increase in payments to inframarginal adopters, while the second term is the subsidy paid to newly induced adopters. Using the elasticity definition, the mechanical spending per induced adoption is

$$\frac{p ds}{dp} = \frac{s}{\varepsilon}. \quad (\text{A.14})$$

Hence total subsidy spending per induced adoption is $s/\varepsilon + s$.

In a standard MVPF, firms that already adopt WFH are policy beneficiaries, since they receive the subsidy as a mechanical transfer. By the envelope theorem, however, firms at the margin of adoption are indifferent between adopting and not adopting, so induced take-up does not generate first-order producer surplus beyond the transfer itself. Hence the total willingness to pay from a small subsidy increase is

$$dWTP = p ds + s dp + EV_H dp, \quad (\text{A.15})$$

where EV_H denotes the household equivalent variation generated by an induced WFH adoption. Net government cost is

$$d\text{Cost} = p ds + s dp - FE dp, \quad (\text{A.16})$$

where FE is the fiscal externality per induced adoption. Dividing by dp , the standard MVPF can be written as

$$MVPF = \frac{s/\varepsilon + s + EV_H}{s/\varepsilon + s - FE}. \quad (\text{A.17})$$

Because this object includes mechanical transfers to firms in the numerator, it is less informative about the effectiveness of subsidies as a tool to reduce the child penalty. For that purpose, we report instead the shadow MVPF

$$MVPF^{\text{shadow}} = \frac{EV_H}{s/\varepsilon + s - FE}, \quad (\text{A.18})$$

which isolates the household surplus created by induced WFH adoption from the transfer received by firms.

The shadow-optimal subsidy equates household benefits to the net fiscal cost of the marginal subsidy increase, implying $MVPF^{shadow} = 1$. Using $s^{\max} = EV_H + FE$, this yields

$$s^{opt} = \frac{\varepsilon}{1 + \varepsilon} s^{\max}. \quad (\text{A.19})$$

This expression highlights that the optimal subsidy scales with the elasticity of adoption, reflecting the trade-off between inducing marginal firms and transferring resources to inframarginal adopters.

B.2 Calibration of fiscal parameters

Tax wedge. We construct an average tax wedge on labor income combining income taxation and social security contributions. Gross annual earnings are obtained by annualizing pre-birth weekly earnings from the analysis sample. Taxable income for the personal income tax (IRPEF) is computed by deducting employee social security contributions. We then apply the statutory 2022 IRPEF schedule (23–35–43 percent brackets) and a stylized schedule of employee tax credits for dependent employment. We exclude regional and municipal surcharges as well as additional tax credits (e.g. *trattamento integrativo*), yielding a conservative estimate of the tax burden.

We compute the net IRPEF liability and express it as a fraction of gross earnings. The total fiscal wedge is then given by the sum of net IRPEF, employee social security contributions (9.19 percent), and employer social security contributions (23.81 percent). This yields an overall wedge of approximately 44 percent, which we apply to the IV estimate of earnings gains to compute the fiscal externality.

Parental leave. We calibrate the fiscal cost of parental leave by combining replacement rates with observed earnings. Weekly gross earnings are obtained from the analysis sample. We consider two replacement rates, 30 and 80 percent, reflecting institutional variation in parental leave generosity. Benefits are computed as a fraction of earnings net of employee contributions and are then adjusted for income taxation using the same effective tax rate as above. The fiscal cost of parental leave includes both the net transfer to the worker and the foregone social security contributions.

Formally, the weekly fiscal cost of leave is given by

$$b_{PL} = B_{PL}^{net} + SSC_{emp} + SSC_{er},$$

where B_{PL}^{net} denotes the after-tax benefit. This implies weekly fiscal costs of approximately 269 euros under the 30 percent replacement rate and 457 euros under the 80 percent replacement rate.

Implementation. The fiscal externality is computed as

$$FE = \tau \Delta Y + b_{PL}(-\Delta PL),$$

where ΔY is measured in annual earnings and ΔPL in weeks. Accordingly, we use weekly parental leave costs and an annual tax wedge. We use the 80 percent calibration as our preferred one.

According to our IV estimates, WFH increases mothers' earnings by about 3,932 euros and reduces parental leave by roughly 5.36 weeks. Under an 80 percent replacement rate for parental leave, we derive a fiscal externality of approximately 4,225 euros, reflecting both higher tax revenues and reduced public spending on leave. Combining these components yields a social surplus of about 6,428 euros per induced contract. This is the maximum subsidy a government would be willing to pay to induce a WFH contract.

This benchmark indicates that the social value of WFH adoption is substantial. However, translating this value into policy requires understanding the responsiveness of firms' adoption decisions. The optimal policy must scale down the subsidy to account for inframarginal transfers and the subsidy paid to newly induced adopters.

C 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% in 2024 (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.