Skatteforsk
Centre for Tax Research

Working Paper No. 14

The Impact of Flexibility at Work on Fertility

June 2024

Bernt Bratsberg
Frisch Centre

Selma Walther
University of Sussex

Skatteforsk's Working Paper Series is financed through support from The Research Council of Norway (Grant number 341289).

School of Economics and Business

Norwegian University of Life Science

The Impact of Flexibility at Work on Fertility

Bernt Bratsberg* Selma Walther†

June 6, 2024

Abstract

Leveraging the first Covid-19 lockdown in Norway as a laboratory for an increase in work flexibility, we uncover a significant and persistent increase in births nine months later. Using the Goldin (2014) measure of work flexibility based on occupation characteristics, we show that fertility increases were concentrated among women in "greedy jobs" with lower flexibility prior to lockdown. We formalise and develop the intuition of Goldin (2014) in a theoretical model where greedy work and greedy children place similar demands on a woman's time. The model explains the mechanism by which an increase in flexibility boosts the fertility of higher earning women, and shows it unfolds under relatively simple theoretical assumptions. The increase in work flexibility under Covid-19 lockdown allowed high-earning women in greedy jobs to alleviate the career-family trade-off.

JEL Classification Code: J13

Keywords: Work flexibility, Fertility, Career-Family Trade-off

*Frisch Centre and Norwegian Institute of Public Health. Email: bernt.bratsberg@frisch.uio.no

†Department of Economics, University of Sussex, IFS and IZA. Email: s.walther@sussex.ac.uk

The authors thank Abi Adams, Alison Andrew, Matthias Doepke and Hanna Wang for useful discussion. This work was supported by the Research Council of Norway through its funding of "Employment, investment, and inequality in the aftermath of the Covid-19 crisis" (project # 315769). Data on loan from Statistics Norway have been essential for the research.

1 Introduction

Women continue to be paid less than men for the same work. In the United States, the gender pay gap was 18.8% in 2010 (Kunze 2018), while in Norway, a context with relatively egalitarian gender norms, the gender pay gap fell to 12.4% in 2022 (Fløtre and Tuv 2023). Recently, theories as to why the gender pay gap persists have centred around temporal flexibility. Goldin (2014), in particular, argues that "greedy work" plays a key role: the ability to work *specific* hours boosts wages. Conversely, women who struggle to provide this temporal flexibility due to the demands of childcare will earn a lower wage.[1] Indeed, flexible working arrangements are associated with a higher wage (Arntz, Yahmed, and Berlingieri 2022) and lower child penalty (Bang 2022).

In this paper, we formally develop Goldin's intuition in a theoretical model of greedy work that also introduces the notion of "greedy kids", and provide empirical evidence that increases in flexibility boost women's fertility in Norway. In the model, our key insight is that children make demands on a woman's time that have parallels with greedy work: the need to look after a sick child, or pick a child up when school ends, also happens at *specific* hours. These hours may coincide with crucial working hours, forcing a woman to make choices between work and family. Childcare provision outside of standard working hours tends to be less developed and more expensive (Henly, Ananat, and Danziger 2006, Bihan and Martin 2004), making it challenging to fully outsource childcare during evening working hours, for example.

In the model, a woman chooses working hours, consumption and number of children to maximise her utility, subject to an endogenous price of child quantity that depends on working hours and the flexibility of the job. The key assumption is that the child price increases with hours worked. This can capture a multitude of interpretations around greedy work: that childcare is more expensive outside of normal working hours, that a woman may feel a disutility from being away from her children for long periods, or, as in Goldin (2014), that missing out on key working hours can result in a wage cost. Crucially, we allow flexibility at work to reduce this child price: for example, working from home can allow a woman to multitask over work and childcare.

We show that, without making any further assumptions, when flexibility increases, the probability of having a child increases for all women. Importantly, fertility increases *more* for women who receive a larger flexibility boost, and those who work more hours and earn more. We take these predictions to the empirical setting.

We leverage the first Covid-19 lockdown in Norway in March 2020 - an unexpected and exogenous event - to study the impact of increases in work flexibility on fertility. Norway experienced distancing measures, travel restrictions and closure of schools and a number of service industries, with most higher-end occupations moving to work from home. First, we document a striking and persistent increase in the number of births, nine months after the first Covid-19 lockdown began. Using a difference-in-difference event study specification with cohort * year fixed effects, we show approximately 0.8 additional monthly births per 1000 women, or 11% of baseline births in the same

[1] These women may also switch to different occupations where the reward to flexibility, and often the overall wage, is lower.

calendar months three years earlier. The fertility increase is concentrated among 25-39 year-old employed women with a partner.

Next, we use occupational characteristics capturing flexibility from Goldin (2014) to categorise women as having low or high flexibility jobs immediately prior to the first lockdown. This categorisation yields a measure of how much lockdown increased flexibility: low flexibility jobs had a larger flexibility increase with lockdown. High flexibility jobs were already flexible and less affected by the move to working from home.

We confirm that the fertility increase was concentrated among women with less flexible jobs prior to lockdown. We interpret this as evidence that the increase in flexibility due to lockdown allowed these women to better reconcile career and family, akin to our theoretical model. We also establish that fertility increased most for women earning above median income before lockdown. Alternative specifications, including changing the control year and introducing placebo lockdowns in other years, confirm the robustness of our findings.

We provide the first evidence that flexibility directly impacts fertility. Existing work highlights the importance of work flexibility for women but has not shown impacts on fertility. For example, women experience more work interruptions during the day (Cubas, Juhn, and Silos 2021), value the option to work from home (Mas and Pallais 2017, Wiswall and Zafar 2017), benefit when fathers' parental leave becomes flexible (Persson and Rossin-Slater, forthcoming) and become more productive with more flexibility (Angelici and Profeta 2020). New to this literature, we show that increases in flexibility boost fertility, and our theoretical model formalises the Goldin (2014) intuition, drawing a new parallel between caring for children and having a greedy job.

Relatedly, flexibility can reduce the gender wage gap. Goldin and Katz (2011) show that a linear wage schedule among pharmacists has helped to shrink the gender wage gap. In a structural framework, allowing women to switch to flexible jobs helps close the gender wage and hours gaps (Hotz, Johansson, and Karimi 2020).

We focus on Covid-19 lockdown as a source of change in flexibility, but other work documents broader changes due to the pandemic. In the United States, the Covid-19 pandemic triggered a baby boom among U.S.-born women, particularly middle-aged women with college degrees (Bailey, Currie, and Schwandt 2023); the latter is similar to our findings for Norway. Similar baby booms were observed in Finland (Nisén et al 2022) and Spain (Cozzani et al 2023). Demographers have emphasised the role that the social security system may have played in the Norwegian fertility increase (Sobotka et al 2023, Lappegård et al 2023); this is consistent with our findings, as the security net reduced the Covid-19 uncertainty Norwegian women faced, allowing changes in flexibility to boost fertility independently.

The paper proceeds as follows. Section 2 outlines a theory of flexible work and fertility. Section 3 outlines the data, shows descriptive statistics and explains the empirical approach. Section 4 presents the main results, with evidence for the impact of flexibility on fertility and a series of robustness checks. Section 5 concludes.

2 A Model of Greedy Work and Greedy Kids

In this section, we present a novel theoretical model of the career-family decision that builds on the intuition in Goldin (2014). Goldin posits that greedy work rewards long hours and particular hours, making it less compatible with family commitments than regular occupations. Our theoretical model draws an interesting parallel between the temporal demands of career and family. The woman has two greedy demands on her time - work and children - and makes optimal choices in light of these.

2.1 Environment

A woman has standard preferences described by $u(c, h, n)$, where c is consumption, h is the number of hours worked, and n is the number of children.

The woman's earnings as a function of working hours are $y(h)$. We allow for potentially nonlinear wages to capture higher marginal compensation for long hours, which Goldin (2014) highlights as a key feature of greedy work. Below, we discuss the implications of this feature for our empirical predictions.

The woman's budget constraint is given by

$$c \leq y(h) - n * z \tag{1}$$

where z measures the price of child quantity. In the classical economic treatment of fertility (e.g., Becker and Lewis 1973, Bhalotra, Venkataramani, and Walther 2023) z is often viewed as being exogenously determined by the market price of child quantity. By contrast, to further explore the consequences of greedy work and greedy children, we allow z to be endogenously determined by the woman's time use.[2] We let z be given by

$$z = q(h, \theta), \tag{2}$$

where θ is a parameter describing the flexibility of the working environment, which we discuss in detail below. We make two further assumptions, which define our key departures from the classical framework.

First, we assume that $\frac{\partial q}{\partial h} > 0$, so that children are more costly per unit for women working long hours. This specification is designed to capture the interplay between greedy work and greedy children. Working long hours in our model is costly not only because of the disutility of labor, but also because of the increasing difficulty of reconciling work with family life when h is large. This can capture a variety of greedy work implications: demanding jobs may require frequent travel or antisocial working hours; alternatively, mothers may struggle with being away from their children for extended periods, or may find the arrangement of ad hoc childcare stressful. Indeed, in our empirical setting in Norway, childcare outside of normal working hours is difficult and costly to

[2] For simplicity of exposition, we abstract away from child quality as a separate good.

arrange.

Second, we assume that $\frac{\partial q}{\partial \theta} \leq 0$ and $\frac{\partial^2 q}{\partial \theta \partial h} \leq 0$, meaning that child-related costs decrease with flexibility and become less sensitive to working hours when work is flexible. For one possible micro-foundation of this assumption, suppose that the woman can work from home and "multitask" for a fraction θ of her total working time, and that child-related costs are driven by the time spent away from home, i.e., $q(h, \theta) = f((1 - \theta)h)$ for an increasing, weakly convex function $f(.)$. It is easy to see that $\frac{\partial^2 q}{\partial \theta \partial h} \leq 0$ in this setting.[3] More broadly, our assumption lets the parameter θ capture the marginal benefits of reconciling work and family, which are more pronounced for women who work long hours.

In this environment, we first characterize and graphically illustrate optimal fertility choices. Then, we derive the key predictions of the model that we take to the data.

2.2 Optimal Fertility Choices

The following lemma, which follows directly from the discussion above, summarizes the woman's optimization problem. To simplify the exposition, we focus on the extensive margin of fertility, setting $n \in \{0, 1\}$.

Lemma 1 *The woman's maximisation problem and her indirect utility, conditional on having n children, is given by*

$$V(n, \theta) = \max_{c,h} \{u(c, h, n) \text{ subject to } c \leq y(h) - n * q(h, \theta)\} \tag{3}$$

For a given set of preferences, wage schedules and flexibility θ, the woman optimally chooses to have a child if and only if $V(1, \theta) \geq V(0, \theta)$.

Figure 1 shows the woman's problem in a simple case where the working day is divided into regular working hours $h \leq h_0$ and additional working hours $h > h_0$. In each panel of the figure, working hours are on the horizontal axis, and consumption is on the vertical axis. For illustration in this figure (but not in our general results) we let wages $y(h)$ and child costs $q(h, \theta)$ be piecewise linear in h. Both marginal wages and marginal child-related costs increase after normal working hours. The thick solid line is the associated budget constraint for $n = 0$ (without a child), and the thick dashed line is the corresponding budget for $n = 1$ (with a child). The thin contours are the woman's indifference curves. Again for illustration only, we assume that the woman's preferences are separable and given by $u(c, h) + bn$, where b stands for the fixed utility benefit of having a child.

Panel (a) shows the optimal labor-consumption tradeoff for a woman who prefers to work only during normal hours of (non-greedy) work. Her optimal choice without children, where her indifference curve is tangent to the solid budget line, yields utility u_0. Her optimal choice with

[3] For an alternative micro-foundation, suppose that an additional fraction $\chi = 1 - \theta$ of working time must be spent on unproductive activities such as commuting, and that child-related costs are driven by time spent away from home, with $q(h, \theta) = f((1 + \chi)h)$ for an increasing function $f(.)$. Again, it is easy to check that $\frac{\partial^2 q}{\partial \theta \partial h} \leq 0$.

children, given by tangency with the dashed budget line, involves slightly less labor supply because of the implicit tax on earnings imposed by child-related costs. This choice yields utility $u_1 + b$. The difference between u_1 and u_0 on the vertical axis is the indirect tax for having a child, and the woman will optimally choose to have a child if and only if her preference for children b is greater than this difference.

Panel (b) shows the corresponding choices for a woman who prefers to work longer hours. As child-related costs are increasing in working hours, this woman faces a larger implicit tax when having a child. Therefore, the difference between u_1 and u_0 is larger for a woman with this preference profile. All else equal, she is less likely to choose to have a child than the woman in panel (a).

Panels (c) and (d) illustrate the changes, for each of the two cases above, when work becomes more flexible. We first discuss the general predictions of the model, which do not depend on the functional forms we draw in the figure, and then discuss these latter illustrations in more detail.

2.3 The Effect of Flexibility on Optimal Fertility

We derive the key empirical predictions of our model, which describe the change in women's optimal choices after an increase in the work flexibility parameter θ. Proofs of propositions are in Appendix A. Applying the envelope theorem to the maximization problem in Equation (3), we can derive the effect of increased flexibility on the woman's indirect utility:

Proposition 1 (Effect of Flexibility on the Value of Children) *The effect of an increase in the work flexibility parameter θ on the woman's maximized utility, conditional on having n children, is:*

$$\frac{\partial V(n;\theta)}{\partial \theta} = n * \lambda(n) * \left(- \frac{\partial q(h,\theta)}{\partial \theta} \bigg|_{h=h^{\star}(n;\theta)} \right) \tag{4}$$

where $\lambda(n)$ is the Lagrange multiplier on the woman's budget constraint, which measures the marginal utility of wealth, and $h^{\star}(n;\theta)$ is her optimal choice of working hours.

This result decomposes the marginal benefits of flexibility into three terms. First, flexibility benefits scale with the number of children n, since costs are assumed to also scale with n, so that the marginal benefit is zero for childless women and stronger for women with more children. Second, it depends on the marginal utility of wealth $\lambda(n)$. Since children are costly, and flexibility leads to a cost saving, mothers with high marginal utility of wealth will experience a stronger utility gain. Third, the effect of flexibility is proportional to the marginal decrease in child costs $q(h,\theta)$, evaluated at the optimal choice of working hours $h = h^{\star}(n;\theta)$.

This characterization leads to the main empirical predictions of our model:

Proposition 2 (Empirical Predictions) *The effects of a marginal increase in the work flexibility parameter θ on women's optimal choices are as follows:*

1. *The incentive to have one child, measured by the distance $V(1;\theta) - V(0,\theta)$, increases.*

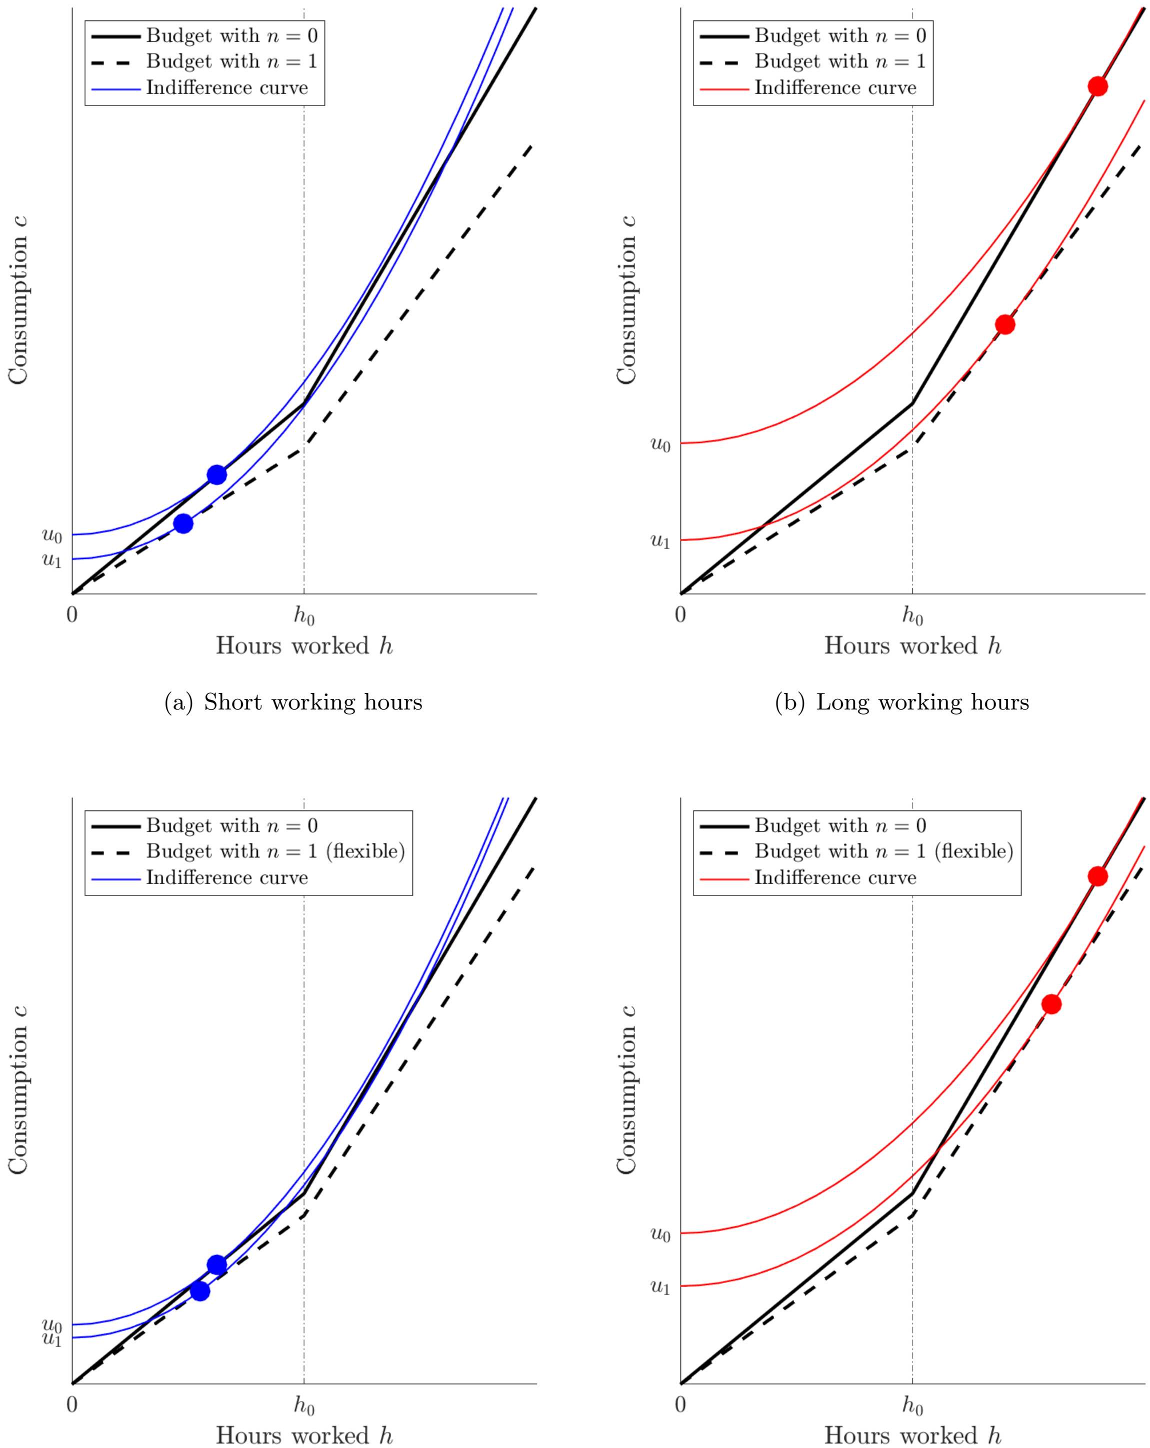

(a) Short working hours

(b) Long working hours

(c) Short hours with flexibility

(d) Long hours with flexibility

Figure 1: **Illustration of Career-Family Trade Off**

2. *The incentive to have one child increases by a greater amount with a greater increase in flexibility θ.*
3. *The incentive to have one child increases by a greater amount for women who work longer hours, conditional on the marginal value of wealth $\lambda(n)$.*

The first two predictions follow directly from Equation (4). It is clear that the value of childlessness $V(0, \theta)$ is unaffected by a change in work flexibility θ. Moreover, the right-hand side is always positive for the value $V(1, \theta)$ of having a child. Therefore, the incentive to have a child always increases in θ. Further, the distance between $V(1; \theta)$ and $V(0, \theta)$ increases with a larger increase in θ, implying that jobs with larger flexibility gains create greater fertility incentives.

The third prediction follows by analyzing how the right-hand side of Equation (4) varies in the optimal choice $h^{\star}(n)$ of working hours. A woman who works longer hours is more affected by a change in flexibility, precisely because child-related costs are more sensitive to flexibility than for women who work shorter hours. In terms of our micro-foundation for $q(h, \theta)$, multitasking is particularly valuable for mothers who work long hours. Women who work long hours tend to have higher income, so this prediction also implies a greater fertility impact for higher earning women conditional on a level of marginal utility of wealth. Of course, if we allow higher earning women to also have a lower marginal utility of wealth, then this effect would be quantitatively smaller.

Our key predictions are illustrated in panels (c) and (d) of Figure 1. In both panels, the budget constraint for women with children reflects a greater degree of flexibility than in panels (a) and (b). Relative to those panels, the implicit tax when having a child is smaller due to the benefits of flexibility. Panel (c) shows the optimal choices of a woman, with flexibility, who prefers to work short hours, with the same preferences as in panel (a). The difference between her utility with and without children, u_1 and u_0, becomes smaller than in panel (a), so that she is more likely to choose to have a child. This demonstrates our first empirical prediction, namely, that the incentive to have children generally increases with flexible work.

Panel (d) shows the optimal choices of a woman who prefers to work long hours, with the same preferences as in panel (b). The difference between u_1 and u_0 also becomes smaller for this woman, meaning she becomes more likely to have children. In addition, it is clear that the *change* in $u_0 - u_1$ due to flexibility is larger for the woman depicted in panel (d), compared to (c). This reflects our third empirical prediction, namely, that the incentive to have children increases more strongly for women who work long hours.

We close by remarking that our key predictions are unchanged when the wage schedule $y(.)$ and child-related costs $q(.)$ are linear in h. However, strict convexity, as drawn in our illustration, makes the predictions more pronounced. On the one hand, if the wage schedule is convex, women are more likely to work long hours and, therefore, be strongly affected by an increase in flexibility. On the other hand, if child-related costs are convex, which is likely if childcare is disproportionately expensive outside school hours, then the differential effect of flexibility on women who work more becomes stronger.[4]

Our empirical predictions derived in this theoretical environment motivate our empirical approach, described in the next section.

3 Empirical Approach

This section describes the data, context and empirical approach for leveraging Covid-19 lockdown in Norway as a change in work flexibility to analyse fertility impacts.

3.1 Norwegian registry data

We collected data from three administrative sources: the central population, the annual income, and the monthly employer-employee ("a-ordningen") registers of Statistics Norway and the Norwegian tax and social insurance administrations. Anonymous personal identifiers allow us to merge records from the three sources and, from the population register, link newborns to their mother. The register identifies the month of birth. For the broad descriptive statistics covering the period 2010-2022, we make no restriction on the population apart from the mother being a registered resident of Norway at the time of birth. For the analysis samples underlying the event and difference-in-difference (DiD) studies described in Section 3.4, we restrict samples to women residing continuously in Norway during the 24-month interval forming the DiD-analysis pre- and post-periods of the lockdown and control cohorts. From the employer-employee register, we retained active job records as of March 12, 2017, and March 12, 2020. For those with multiple jobs, we kept the record with highest pay. From these records, we extracted the 4-digit ISCO-88 occupation code, which we used to construct work flexibility indices discussed in Section 4.2.

3.2 Covid-19 and lockdown in Norway

The Covid-19 pandemic hit the Norwegian labor market on March 12, 2020, with the unexpected announcement by the Prime Minister that there would be lockdowns and strict regulations on social distancing. There were travel restrictions, and many service industries closed. Most high-end jobs moved to work from home. The lockdown policy ended in March 2021.

3.3 Summary statistics

Figure 2 shows total monthly births (Panel A) and births per 1000 females (Panel B) between January 2010 and December 2022 in Norway. Births are highly seasonal and usually peak in the summer months. Further, fertility declined over this period, similar to trends in other high income countries.

[4] In addition, strict convexity makes the predictions more robust to alternative assumptions. For example, in our leading micro-foundation, if we assumed that a fixed number of hours can be spent multitasking (as opposed to a fixed fraction of h), then the third prediction would require strict convexity of $q(.)$.

In Panel B, a divergence from birth trends is clear nine months after the beginning of the first Covid-19 lockdown (indicated by the vertical red dashed line), particularly among the 25-39 age group, whose fertility during this period is higher than in the previous year. Fertility stops rising for this group approximately a year later, coinciding with the end of lockdowns in Norway and a gradual return to work in person.

Figure 2: Births in Norway over time

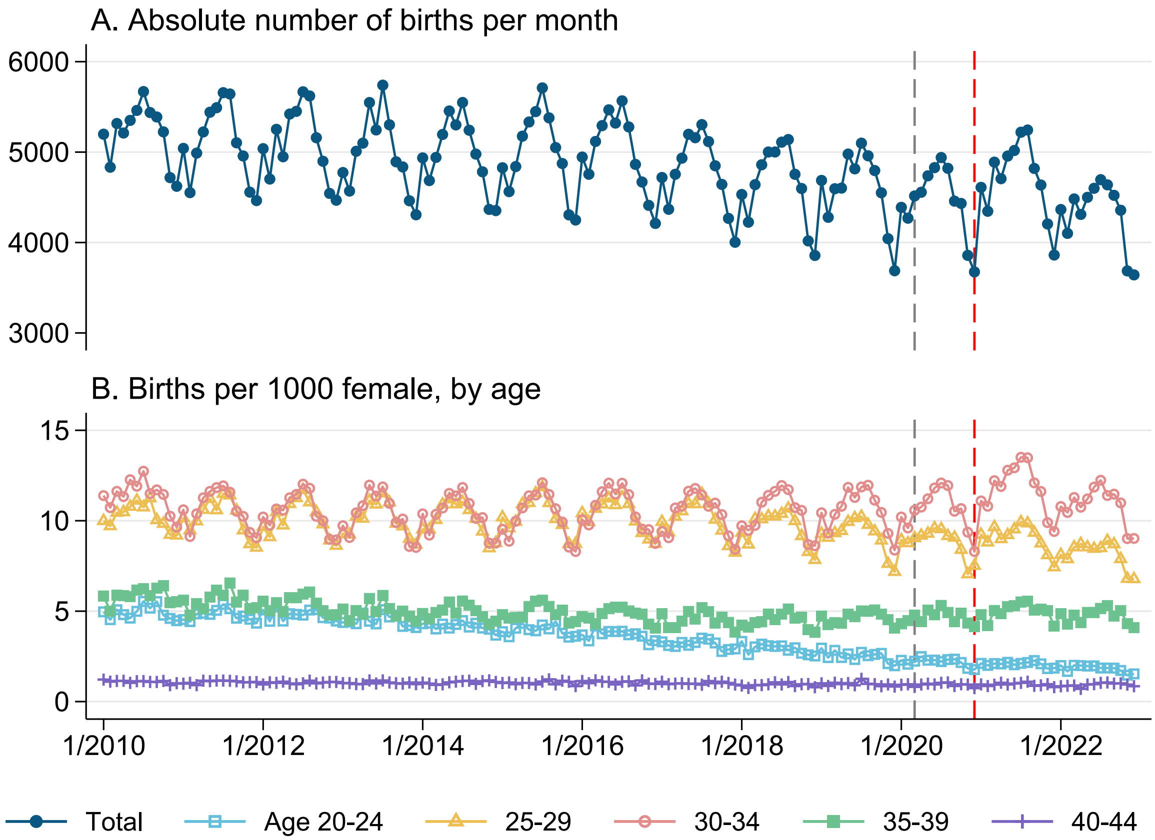

Notes: The grey dashed line indicates the start of the first Covid-19 lockdown, and the red dashed line is placed nine months later, when the first births of lockdown conceptions would have occurred. There were 753,965 births during this period.

3.4 Empirical Approach

The descriptive evidence showing an increase in fertility during lockdown could be driven by other trends in the same period. To verify the role of lockdown, we take a difference-in-difference event study approach using individual-level birth records from the Norwegian registry data. We use individual birth records, in combination with labour market records, to construct a monthly panel dataset of Norwegian women. This panel dataset includes all women, not only those with births.

We compare monthly birth probabilities of these women in a 36-month window around the start of the first lockdown to births in the same range in an earlier year. A child conceived in March 2020, the first month of lockdown, would have been born in December 2020. Therefore, we center the data around December 2020, when the first lockdown conceptions would have been born, and analyse births 12 months prior and 24 months after. We compare these births to a control window around December 2017. By aligning the calendar months, we remove any noise from the seasonality of births.

We choose 2017 to center the control event because it was a relatively unremarkable year in demographic terms, and because a 24-month window after December 2017 brings us to December 2019, shortly before Covid-19 appeared. Still, as a robustness check, we show in Section 4.3 that changing the control year does not change our results.

Given the descriptive findings in Figure 2 that show most births happen during the age range 25-39, we focus our empirical analysis on this group. If a woman is observed in the data to have given birth, we set her subsequent nine observation months to missing, so that an additional birth can occur earliest nine months after the observed birth.[5]

The estimating equation is:

$$z_{i,t,\tau} = \sum_{\tau=-12}^{\tau=+24} \alpha_\tau Month_{i,\tau} + \sum_{\tau=-12}^{\tau=+24} \beta_\tau Covid_{i,t} * Month_{i,\tau} + \gamma Cohort_i * Covid_{i,t} + \eta_{i,t,\tau}, \qquad (5)$$

where $z_{i,t,\tau}$ is the birth outcome for individual i, τ defines observation months and t is observation year (2017 or 2020). $Month_{i,\tau}$ is a dummy variable representing calendar months, centred around December. $Covid_{i,t}$ indicates whether the year is 2020 or 2017. The coefficient β_τ gives the differential impact of Covid-19 lockdown on births, compared to the 2017 birth probability trajectories captured in $\alpha_\tau Month_{i,t}$. We include a full set of cohort * year fixed effects ($Cohort_i * Covid_{i,t}$). Standard errors are clustered at the individual level. Identification relies on 2017 birth probabilities providing a valid counterfactual trajectory for the outcomes of individuals potentially giving birth in 2020, had there been no Covid-19 lockdown, and after allowing for time-varying cohort effects through cohort * year fixed effects.[6]

Our extended time period of analysis, 12 months before and 24 months after the first post-Covid births, allows us to closely monitor the evolution of outcomes before the first lockdown and check that birth trends evolved in a similar way prior to December 2017 and 2020. In Section 4.3, we introduce placebo lockdowns in 2016 and 2018 and show null effects on fertility, confirming the reliability of our findings. To aid our understanding of effect sizes, and check whether average birth probabilities were significantly different before compared to after lockdown, we also estimate

[5] This assumes a woman can conceive again immediately after giving birth, yielding an upper bound on the number of potential births in any given period.

[6] We omit individual fixed effects because they are computationally demanding to estimate given the size of our dataset, but we have verified that the baseline estimates are broadly unchanged with the inclusion of individual fixed effects.

a binary difference-in-difference version of Equation (5) where monthly dummies are replaced with post-December dummies.[7] These regressions yield differences in average birth probabilities in a symmetric window 12 months before and 12 months after December 2020, compared to December 2017 (Appendix Figure B.1, and discussed throughout the text).

4 Results

4.1 Overall effect of lockdown on fertility

Figure 3 shows the results from estimating Equation (5): it is the treatment effect of Covid-19 lockdown on births among women aged 25-39, relative to baseline births in the same months in the control years, conditioning on cohort * year fixed effects.

There is a significant and persistent increase in births, nine months after lockdown began and onwards.[8] The effect is approximately +0.8 additional births per 1000, relative to a baseline mean of 7.3, or 11% of baseline. This is consistent with the first prediction in Proposition 2: an overall increase in flexibility boosts fertility probabilities for all women. The effect persists for twelve months from the first lockdown month with significantly higher births until December 2021, implying significantly higher conceptions until March 2021; importantly, this is when lockdown policies ended. In total, we estimate around 4,300 additional births during this period. That the fertility effect is already seen in the first month is not surprising given that conceptions are most likely to occur in the first month of trying, with a probability of around 30% (Taylor 2003).

The extended time window of analysis allows us to verify our assumption that birth trends in 2017 and 2020 were similar. Prior to the baseline month of December, all coefficients except one are statistically insignificant, indicating no monthly differential pre-December birth trends between 2017 and 2020.

4.2 Mechanisms and the role of flexibility at work

With the advent of the first Covid-19 lockdown, suddenly and without anticipation, the nature of work changed and many jobs moved into homes. Holgersen, Jia, and Svenkerud (2021) show that 38% of all Norwegian jobs could be performed at home during lockdown, but that the figure is substantially higher for more educated professions, e.g. 65.7% for managers but only 7% for machine operators. They also show that the highest paid occupations had the highest rate of working from home.

[7] Specifically, we restructure the data to individual women's birth outcomes in the 12 months before and 12 months after December 2017 and 2020, and estimate:

$$z_{i,t,\tau} = a * PostDecember_{i,\tau} + b * PostDecember_{i,\tau} * Covid_{i,t} + g * Cohort_i * Covid_{i,t} + \eta_{i,t,\tau}. \quad (6)$$

To obtain effects for groups, such as employed vs not employed women, $PostDecember$ and $PostDecember * Covid$ are interacted with these group dummies.

[8] Precisely, monthly data means the first estimated coefficient is 9.5 months after day one of lockdown.

Figure 3: The effect of Covid-19 lockdown on births in Norway for women, ages 25-39

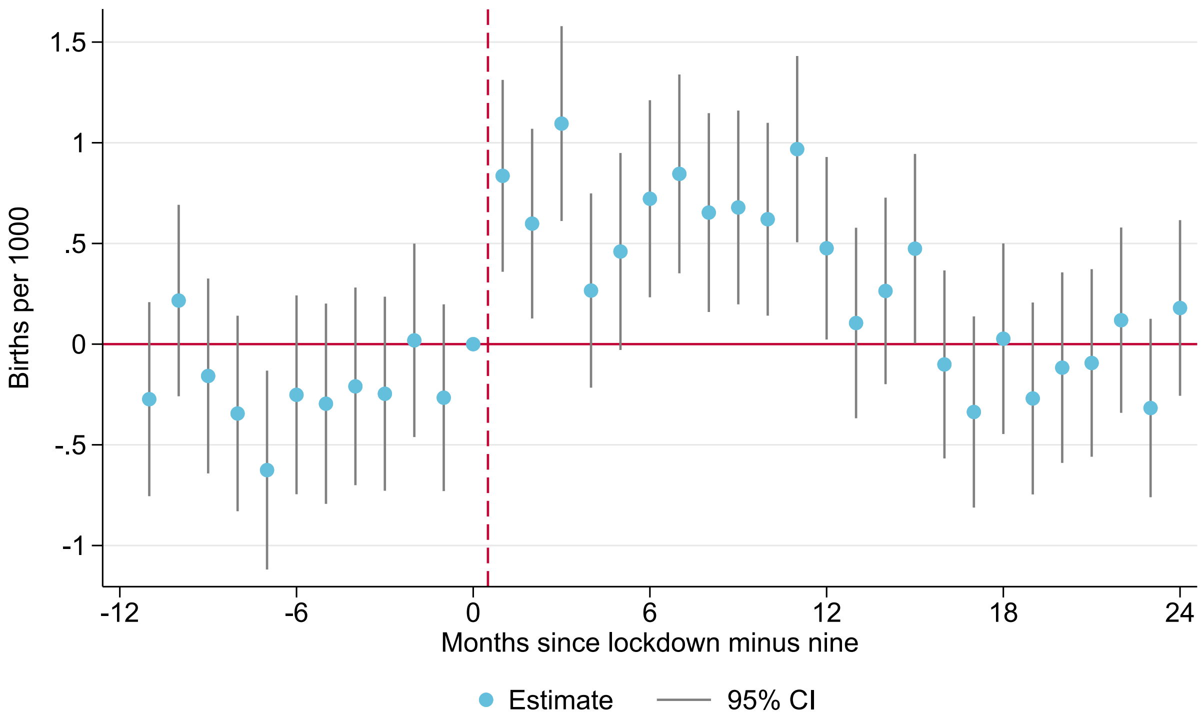

Notes: Treatment cohort consists of women present in Norway between Jan 1, 2020, and Dec 31, 2021, and who were 25 to 39 years of age Dec 31, 2020; control cohort analogously defined for women present in Norway from Jan 1, 2017. Estimation sample has 34,689,929 observations. See Equation (5) for estimation model.

We argue that the rise in flexibility during lockdown was the key driver that stimulated aggregate fertility during this period, as in our theoretical model. We begin by demonstrating that fertility increases were only evident among women with jobs and partners prior to lockdown - necessary conditions for fertility to respond to changes in the nature of work during lockdown. Then, we provide direct evidence for the predictions of our model by showing that fertility increased more for women with larger increases in flexibility, and those with higher incomes.

First, we show that fertility effects were concentrated among working women with an existing partner. Panels A and B of Figure 4 compare the fertility response of women employed and unemployed as of March 12, 2020 (day one of lockdown) relative to the control group of employed or unemployed women on March 12, 2017. We find that the aggregate fertility response was entirely driven by women with a job, with +1 additional births per 1000, or 12 percent of baseline. Moreover, for the unemployed group, coefficient estimates reveal imbalance in pregnancies in the pre-period, implying impacts for unemployed women are difficult to interpret.

Lockdown placed restrictions on outdoor movement and social meetings, which made it challenging to meet partners for those who were single. Figure 4 confirms the intuition that among the employed women seen in Panel A, those in couples saw a fertility boost (Panel D), while single women did not (Panel C). Here, couples are defined as married or cohabiting individuals. Indeed, the average fertility impact among coupled, working women was +1.1 additional monthly births per 1000, or 11 percent of baseline. This is 1.4 times the size of the overall population impact, indicating that coupled, working women aged 25-39 constituted the majority of the fertility response to Covid-19 lockdown.

Figure 4: The effect of Covid-19 lockdown on births in Norway, by women's baseline partner and employment status

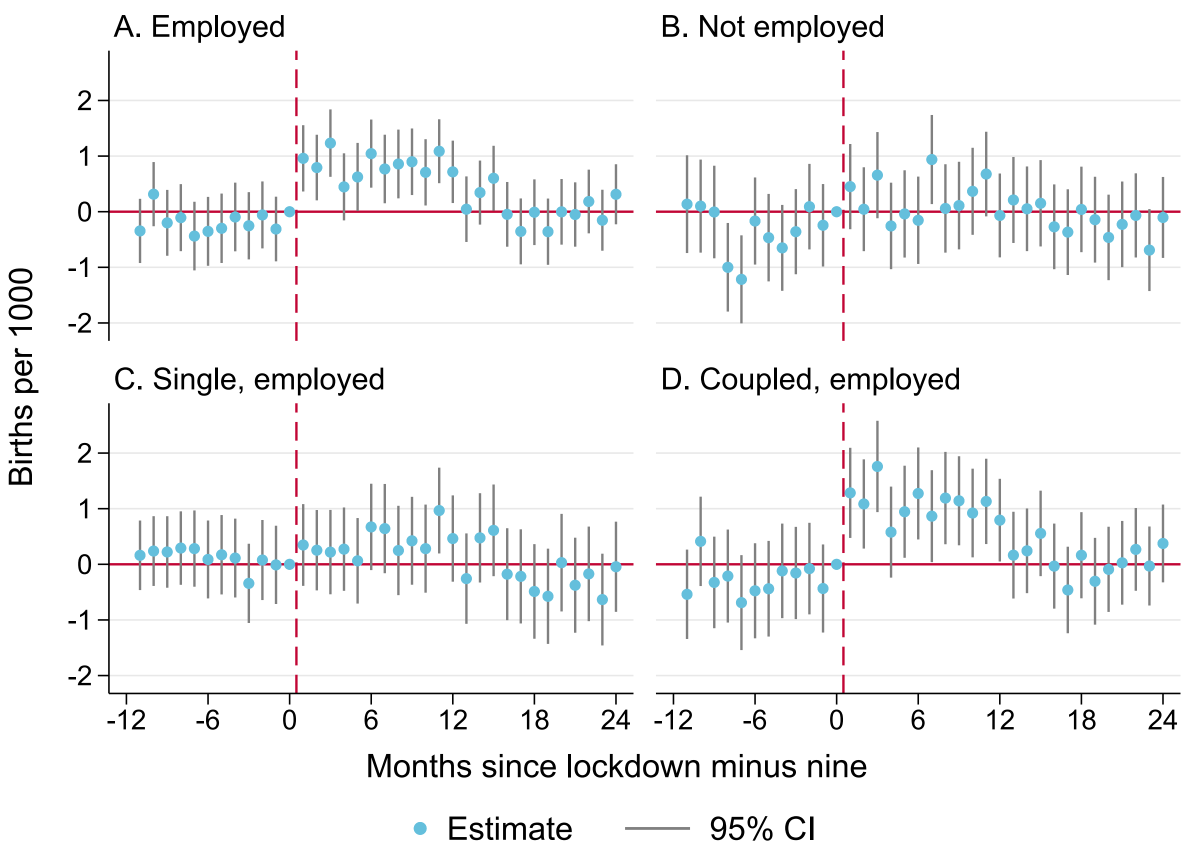

Notes: Employment status measured March 12, 2017 (Control) or 2020 (Treatment), and marital status January 1 of the same year. Being in a couple includes being legally married as well as cohabiting with a partner. Panels A and B are estimates from a single regression of Equation (5), where $Month_{i,\tau}$ dummies are fully interacted with employment status. A similar approach in Panels C and D restricts the sample to employed women and interacts $Month_{i,\tau}$ with partner status. Panels A and B have a sample size of 34,689,929, of which 70.3% of the sample is employed. Panels C and D have a sample size of 24,399,340, of which 67.7% are partnered.

Our theoretical model predicts a greater fertility response among women experiencing a greater flexibility increase. Next, we demonstrate that among coupled, working women, those with the largest increase in flexibility had the largest fertility response. Goldin (2014) argues that five characteristics are associated with a worker having fewer substitutes, and therefore lower flexibility: high time pressure, high contact with others, high maintenance of interpersonal relationships, structured work, and freedom to make decisions. We construct a measure of work flexibility for each woman's occupation on day one of lockdown or control lockdown (12th March 2020 or 2017), using these characteristics extracted from the O*Net data. We split the sample along the median of work flexibility.

Consider two women, one with a low flexibility and one with a high flexibility job at the start of lockdown. The common Covid-19 lockdown shock would have led to a larger absolute increase in flexibility for the low flexibility woman, relative to the high flexibility woman. Consequently, our model predicts a larger fertility increase for the low flexibility woman. We rely on this approach, where baseline levels of a variable are used to predict absolute changes and therefore the intensity of a treatment (as in Acemoglu and Johnson (2007), Bhalotra, Venkataramani, and Walther (2023) and Ager, Hansen, and Jensen (2017)), to provide evidence that flexibility changes were the key driver of fertility increases.

The results in Panel B of Figure 5 point to fertility increasing more for women with low flexibility prior to lockdown. We do not need to rule out positive responses among high flexibility women (Panel A), as they may have also experienced a (smaller) increase in flexibility. Importantly, their responses are smaller in magnitude, and significant in fewer months, compared to women with lower levels of pre-lockdown flexibility (Panel B). A pre-post difference-in-difference comparison of birth outcomes 12 months before and 12 months after December shows that low flexibility women had a significant, 43% larger birth response than high flexibility women (Appendix Figure B.1).

Our theoretical model also predicts a greater fertility response among women working longer hours. As the registry data does not provide a precise record of working hours, we use income as a proxy for hours. We calculate women's earnings rank based on their three best pre-January (2017 or 2020) earnings years, and split women into above and below median earnings rank. Panel B shows that women with above median earnings rank had increases in birth rates in the first three months that were twice as large as those seen among all coupled, employed women. The pre-post difference-in-difference specification shows that high income women had a fertility response almost double that of low income women (Appendix Figure B.1).

Together, these results are consistent with the predictions of Proposition 2: when flexibility increases, women experiencing a larger flexibility shock, or higher earning women, have a larger increase in fertility. In the next section, we explore other mechanisms and show that an increase in work flexibility is the only robust explanation for the fertility effects.

Figure 5: The effect of Covid-19 lockdown on births in Norway, by baseline earnings rank and flexibility of work

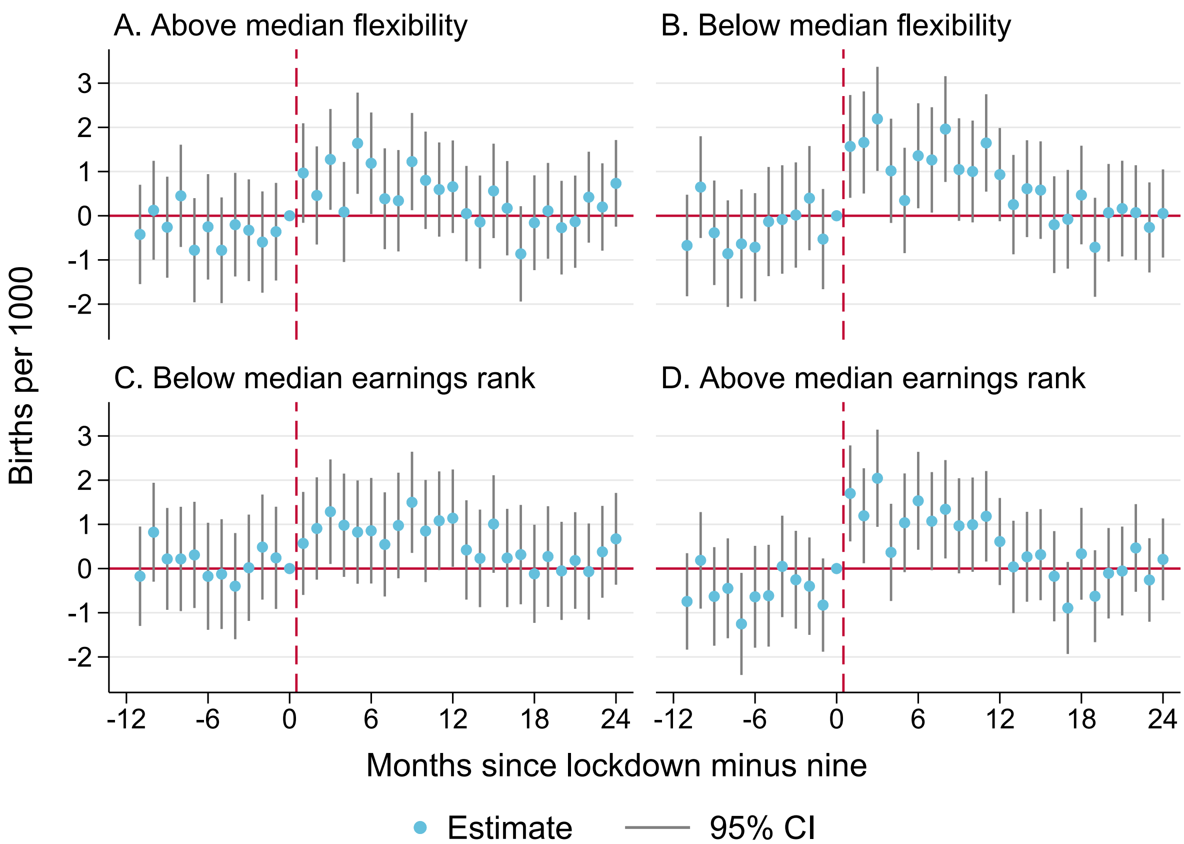

Notes: Occupational flexibility measured on first day of lockdown or control lockdown using the Goldin (2014) work flexibility scale. Earnings rank, within birth cohort, is based on the best three of ten prior earnings years (see text). Panels A and B are estimates from a single regression of Equation (5) where $Month_{i,\tau}$ dummies are fully interacted with a dummy for having below median (low) or above median (high) work flexibility; the sample are coupled, working women and the size is 16,476,823. Similarly, Panels C and D are from a parallel regression where monthly dummies are fully interacted with a dummy for having below median or above median earnings rank; the sample size is 16,507,397.

4.3 Other mechanisms and robustness checks

In our main results, we observe larger increases in fertility among high earning, coupled women in Norway and those working in inflexible jobs pre-lockdown. In this section, we show that this is not driven by changes in the opportunity cost of time, or by the choice of control year, and that a "placebo lockdown" in other years shows null fertility effects.

Job loss or income decline could change the opportunity cost of time. However, children are thought to be a normal good, albeit with a smaller income elasticity compared to the quality of those children (Becker 1960, Doepke 2015). At the same time, fertility tends to be procyclical (Chatterjee and Vogl 2018), declining during uncertain times (Schaller, Fishback, and Marquardt

2020). The Norwegian economy shrank by 4.3% in 2020, relative to 2019 (Statistics Norway), and the furlough scheme replaced around two thirds of income for those jobs that could not be continued under lockdown, implying an income decline for some women. In addition, previous work has shown that female job loss has a negative effect on fertility (Huttunen and Kellokumpu 2016, Currie and Schwandt 2014, Bono, Winter-Ebmer, and Weber 2012).

To evidence that income or job loss is not a competing mechanism, we use our main estimation sample of employed, coupled women and compare fertility responses during lockdown between those that experienced at least one month without pay during lockdown, to those that had pay in all months. We find that the fertility response was not significantly different between these two groups, negating a role for the opportunity cost channel (results not reported for compactness).

Our findings are robust to choice of control year. In Panels A and B of Appendix Figure B.2, we show the estimated coefficients from Equation (5) using 2016 or 2018 as control years, respectively. Our main findings are unchanged: in both cases, effects hover around +1 birth per 1000, similar to the effect size in Figure 3.

Finally, we introduce a placebo lockdown in 2016 and 2018. This specification keeps 2017 as the control year, but assumes a lockdown started in March 2016 or 2018. Panels C and D of Appendix Figure B.2 show null effects in both instances. This confirms our results are not driven by other trends or occurrences during this period.

5 Conclusion

While many of the structural drivers of the gender wage gap, including education and experience gaps, have disappeared, women continue to earn less than men for the same work. Recent research has argued that high-paying careers make specific temporal demands that are difficult to combine with family life.

Our contribution is two-fold. First, we derive a model that formalises the intuition of greedy work and draws the parallel that children place similar needs on a woman's time. It predicts that when flexibility at work increases, birth probabilities increase for all women, but more for women who experience a larger flexibility boost, or who work long hours and earn more. Second, we provide novel empirical evidence that changes in work flexibility drive fertility. Using the first Covid-19 lockdown in Norway, we show that the probability of giving birth increased most for women with less flexible jobs prior to lockdown, and above median earnings, consistent with the theoretical model.

Our findings do not preclude declines in the child penalty as an additional mechanism (Kleven, Landais, and Soegaard 2019). If an increase in work flexibility "levels the playing field" between mothers and non-mothers, this should have an additional stimulating effect on fertility. Until now, discussion of declining fertility has focused on policies such as maternity leave and childcare provision (Doepke, Hannusch, Kindermann, and Tertilt 2023). Our findings point to the importance of another dimension - flexibility at work - that has the power to drive fertility decisions, and may become increasingly important as the nature of work changes.

References

Acemoglu, D. and S. Johnson (2007). Disease and development: The effect of life expectancy on economic growth. *Journal of Political Economy 115*(6), 925–985.

Ager, P., C. W. Hansen, and P. S. Jensen (2017). Fertility and Early-Life Mortality: Evidence from Smallpox Vaccination in Sweden. *Journal of the European Economic Association 16*(2), 487–521.

Angelici, M. and P. Profeta (2020). Smart-working: Work flexibility without constraints. *CESifo Working Paper* (8165).

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

Bailey, M. J., J. Currie, and H. Schwandt (2023). The Covid-19 Baby Bump in the United States. *Proceedings of the National Academy of Sciences 120*(34).

Bang, M. (2022). Job Flexibility and Household Labor Supply: Understanding Gender Gaps and the Child Wage Penalty. *Mimeo*.

Becker, G. S. (1960). *An Economic Analysis of Fertility*. Princeton University Press.

Becker, G. S. and H. G. Lewis (1973). On the interaction between the quantity and quality of children. *Journal of Political Economy 81*(2, Part 2), S279–S288.

Bhalotra, S., A. Venkataramani, and S. Walther (2023). Women's fertility and labor market responses to a health innovation. *Journal of the European Economic Association* (21.4).

Bihan, B. L. and C. Martin (2004). Atypical working hours: Consequences for childcare arrangements. *Social Policy & Administration 38*(6), 565–590.

Bono, E. D., R. Winter-Ebmer, and A. Weber (2012). Clash of career and family: Fertility decisions after job displacement. *Journal of the European Economic Association 10*(4), 659–683.

Chatterjee, S. and T. Vogl (2018). Escaping Malthus: Economic Growth and Fertility Change in the Developing World. *The American Economic Review 108*(6), 1440–1467.

Cozzani, M., P. Fallesen, G. Passaretta, J. Härkönen, and F. Bernardi (2023). The consequences of the Covid-19 pandemic for fertility and birth outcomes: Evidence from Spanish birth registers. *Population and Development Review*, 1–24.

Cubas, G., C. Juhn, and P. Silos (2021, May). Work-care balance over the day and the gender wage gap. *AEA Papers and Proceedings 111*, 149–53.

Currie, J. and H. Schwandt (2014). Short- and long-term effects of unemployment on fertility. *PNAS 111*(41), 14734–14739.

Doepke, M. (2015). Gary Becker on the quantity and quality of children. *Journal of Demographic Economics 81*(1), 59–66.

Doepke, M., A. Hannusch, F. Kindermann, and M. Tertilt (2023). Chapter 4 - The economics of fertility: A new era. In S. Lundberg and A. Voena (Eds.), *Handbook of the Economics of the Family, Volume 1*, Volume 1 of *Handbook of the Economics of the Family*, pp. 151–254. North-Holland.

Fløtre, I. A. and N. Tuv (2023, March). Slik kan lønnsforskjellen mellom kvinner og menn forklares. Technical report, Statistics Norway.

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

Goldin, C. and L. F. Katz (2011). The cost of workplace flexibility for high-powered professionals. *The Annals of the American Academy of Political and Social Science 638*, 45–67.

Henly, J. R., E. O. Ananat, and S. K. Danziger (2006). Nonstandard work schedules, child care subsidies, and child care arrangements. *Mimeo*.

Holgersen, H., Z. Jia, and S. Svenkerud (2021). Who and how many can work from home? Evidence from task descriptions. *Journal of Labor Market Research 55*(4).

Hotz, J. V., P. Johansson, and A. Karimi (2020). Parenthood, family friendly workplaces, and the gender gaps in early work careers. *NBER Working Paper* (24173).

Huttunen, K. and J. Kellokumpu (2016). The effect of job displacement on couples' fertility decisions. *Journal of Labor Economics 34*(2), 403–442.

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

Kunze, A. (2018). The Gender Wage Gap in Developed Countries. In L. M. A. Averett, Susan L. and S. D. Hoffman (Eds.), *The Oxford Handbook of Women and the Economy*.

Lappegård, T., T. Kornstad, L. Dommermuth, and A. P. Kristensen. Understanding the positive effects of the COVID-19 pandemic on women's fertility in Norway. *Population and Development Review*.

Mas, A. and A. Pallais (2017, December). Valuing alternative work arrangements. *American Economic Review 107*(12), 3722–59.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995). *Microeconomic theory*, Volume 1. Oxford University Press.

Nisén, J., M. Jalovaara, A. Rotkirch, and M. Gissler (2022). Fertility recovery despite the COVID-19 pandemic in Finland? *Finnish Journal of Social Research 15*, 25–44.

Persson, P. and M. Rossin-Slater (Forthcoming). When dad can stay home: Fathers' workplace flexibility and maternal health. *American Economic Journal: Applied Economics*.

Schaller, J., P. Fishback, and K. Marquardt (2020, May). Local Economic Conditions and Fertility from the Great Depression through the Great Recession. *AEA Papers and Proceedings 110*, 236–40.

Sobotka, T., V. Skirbekk, and D. Philipov (2011). Economic recession and fertility in the developed world. *Population and Development Review 37*(2), 267–306.

Taylor, A. (2003). ABC of Subfertility: Extent of the problem. *British Medical Journal 327*(7412).

Wiswall, M. and B. Zafar (2017, 08). Preference for the Workplace, Investment in Human Capital, and Gender. *The Quarterly Journal of Economics 133*(1), 457–507.

Online Appendix

A Proofs

A.1 Proof of Proposition 1

Recall the woman's maximization problem from Lemma 1:

$$V(n,\theta) = \max_{c,h} \{u(c,h,n) \text{ subject to } c \leq y(h) - n * q(h,\theta)\}$$

The woman's Lagrangian in this maximization problem, given fixed values of n and θ, is

$$\mathcal{L}(n;\theta) = u(c,h,n) - \lambda[c - y(h) + n * q(h,\theta)]$$

Let $\lambda(n)$ be the Lagrange multiplier associated with the maximizing choices $c^\star(n;\theta)$ and $h^\star(n;\theta)$. The Envelope Theorem (see Mas-Colell, Whinston, and Green 1995, Theorem M.L.1) directly implies that

$$\begin{aligned}\frac{\partial V(n;\theta)}{\partial \theta} &= \frac{d}{d\theta}\left[u(c^\star(n;\theta), h^\star(n;\theta), n)\right] \\ &= -\lambda(n) n \frac{\partial q(h^\star(n;\theta),\theta)}{\partial \theta}\end{aligned}$$

which establishes the result in the proposition.

A.2 Proof of Proposition 2

Using Proposition 1, and using the fact that $\frac{\partial V(0;\theta)}{\partial \theta} = 0$, we obtain

$$\begin{aligned}\frac{\partial\left(V(1;\theta) - V(0;\theta)\right)}{\partial \theta} &= \frac{\partial V(1;\theta)}{\partial \theta} \\ &= -\lambda(1)\frac{\partial q(h^\star(1;\theta),\theta)}{\partial \theta}\end{aligned}$$

We use this equation to establish the three claims in the proposition. First, since we have assumed that $\frac{\partial q(h,\theta)}{\partial \theta} \leq 0$, the expression above is always positive, which establishes point 1 in the proposition.

Second, consider a discrete increase in flexibility from an initial value θ_0 to a larger value $\theta_0 + \delta$, with $\delta > 0$. The incentive to have one child changes by

$$\begin{aligned}
[V(1;\theta_0+\delta) - V(0;\theta_0+\delta)] - [V(1;\theta_0) - V(0;\theta_0)] &= \int_{\theta_0}^{\theta_0+\delta} \left[\frac{\partial \left(V(1;\theta) - V(0;\theta) \right)}{\partial \theta} \right] d\theta \\
&= \int_{\theta_0}^{\theta_0+\delta} \left[-\lambda(1;\theta) \frac{\partial q(h^\star(1;\theta),\theta)}{\partial \theta} \right] d\theta
\end{aligned}$$

where we have made explicit the dependence of the marginal value of wealth on θ. Since every term under the integral is positive, the total change in incentives is increasing in the change in flexibility δ, which establishes point 2 in the proposition.

Third, consider two women denoted A and B who have the same marginal value of wealth $\lambda(1)$ conditional on having a child, and the same initial flexibility θ. Assume that woman A (due to different preference/disutility of labor, for example), initially chooses longer working hours, with $h^\star_A(n;\theta) > h^\star_B(n;\theta)$ for all n. Then the change in incentives to have a child for woman A, after an increase in θ, is

$$\begin{aligned}
\frac{\partial \left(V_A(1;\theta) - V_A(0,\theta) \right)}{\partial \theta} &= -\lambda(1) \frac{\partial q(h^\star_A(1;\theta),\theta)}{\partial \theta} \\
&= -\lambda(1) \left[\frac{\partial q(h^\star_B(1;\theta),\theta)}{\partial \theta} + \int_{h^\star_B(1;\theta)}^{h^\star_A(1;\theta)} \left[\frac{\partial^2 q(h,\theta)}{\partial \theta \partial h} \right] dh \right] \\
&\geq -\lambda(1) \left[\frac{\partial q(h^\star_B(1;\theta),\theta)}{\partial \theta} \right] \\
&= \frac{\partial \left(V_B(1;\theta) - V_B(0,\theta) \right)}{\partial \theta}
\end{aligned}$$

where the inequality follows from our assumption that $\frac{\partial^2 q(h,\theta)}{\partial \theta \partial h} \leq 0$. Hence the change in incentives to have a child in response to an increase in θ is larger for woman A than for woman B, which establishes point 3 in the proposition.

B Additional tables and figures

Figure B.1: The effect of Covid-19 lockdown on Births per 1000 women in Norway: a comparison of average birth outcomes in the 12 months before and 12 months after December 2017 and 2020

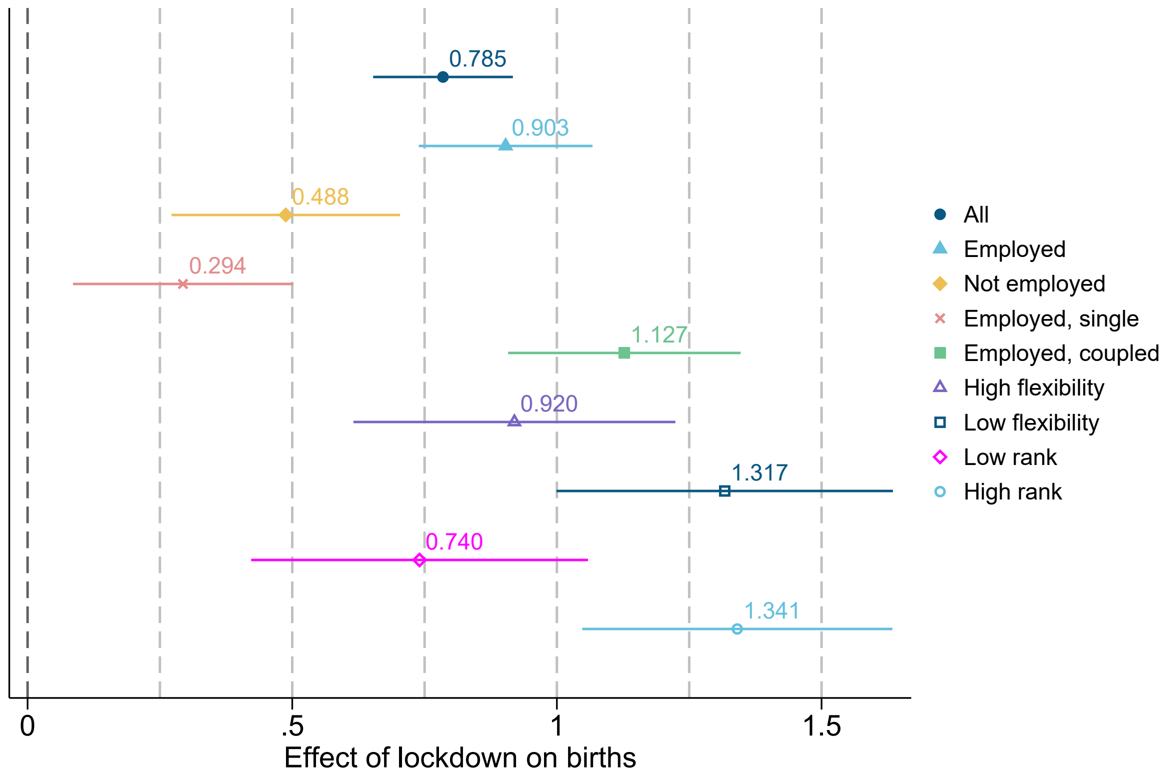

Notes: This figure shows estimates from a regression of Equation (6), described in Footnote 7, where monthly event study dummies are replaced with a single pre- and post-December dummy, interacted with treatment: $PostDecember * Covid$. To obtain effects for various groups (e.g. employed vs not employed), we interact $PostDecember * Covid$ with the appropriate group dummies (e.g. employed dummy). The coefficients therefore show the difference in average birth probabilities in the 12 months before and 12 months after December 2020, compared to the same difference around December 2017. The sample size for the estimate for *All* women is 2,064,452. Estimates for *Employed* and *Not employed* are from one regression with a sample size of 2,064,452 and the p-value for the test of equality of the coefficients on dummies employed vs. not employed interacted with $PostDecember * Covid$ is 0.003. *Employed, single* and *Employed, coupled* coefficients are from one regression of 1,458,928 employed women; the p-value for equality of coefficients on the partner status dummies interacted with $PostDecember * Covid$ is 0.000. Estimates for *High flexibility* and *Low flexibility* are from one regression of 1,002,384 employed, coupled women, where the p-value for equality of coefficients on the flexibility dummies interacted with $PostDecember * Covid$ is 0.077. Coefficients on *Low rank* and *High rank* are from one regression of 1,004,252 employed, coupled women, where the p-value for equality of coefficients on the earnings rank dummies interacted with $PostDecember * Covid$ is 0.007.

Figure B.2: Alternative control years and placebo lockdowns

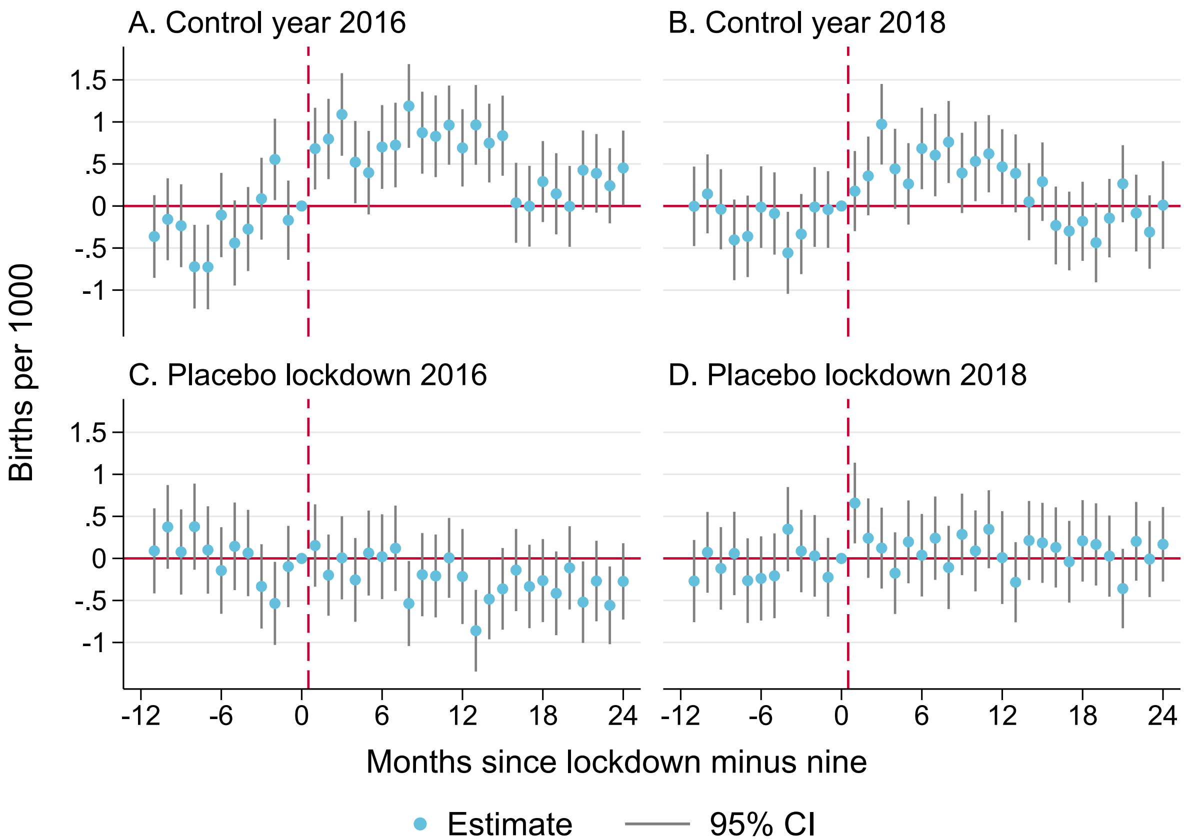

Notes: These figures show estimated coefficients on monthly dummies from Equation (5), estimated for all women as in Figure 3. Panels A and B show estimates with alternative control lockdowns in March 2016 and 2018, where 2020 is kept as the treatment year; sample sizes are 34,410,713 and 34,951,770 respectively. For the latter regression, March-November 2020 will be in the control lockdown period, so that the last nine coefficients in Panel B may be subject to noise from a change in the timing of births (e.g. premature births) during the control period. However, estimated effects for the first twelve post months are not subject to this comment. Panels C and D show results using placebo lockdowns in March 2016 and 2018, where the control lockdown is kept at March 2017. Due to our 36-month window, there is some overlap between the control and treatment months outside of the 6 months before and 6 months after December of the treatment and control years. Sample sizes are 33,814,798 and 34,355,855 respectively.