Greedy Work and Greedy Kids: How Flexibility at Work Shapes Fertility

Bernt Bratsberg* Selma Walther†

April 2, 2024

Abstract

Using 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 kids 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 Codes: J12, J13

Keywords: Greedy work, 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

1 Introduction

Women continue to be paid less than men for the same work. In Norway, a context with relatively egalitarian gender norms, the gender pay gap has shrunk from 16.1% to 12.4% between 2002 and 2022 (Fløtre and Tuv 2023). Most recently, theories as to why the gender pay gap persists have centred around the notion of temporal flexibility. Goldin (2014), in particular, argues that "greedy work" is responsible: the number of working hours drives wages, but also *specific* hours. These demands on a woman's time make it challenging to reconcile career with children, another temporally specific aspect of life.

In this paper, we develop Goldin's intuition more formally in a theoretical model that introduces the notion of "greedy kids", and provide empirical evidence for the importance of flexibility in driving 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 *particular* hours. These hours may coincide with crucial working hours, forcing a woman to make choices between work and family. Both publicly provided and private formal childcare are less developed outside of standard working hours, with private or ad-hoc arrangements being the norm, and prices for paid care tending to be higher (Henly, Ananat, and Danziger 2006, Bihan and Martin 2004, Verhoef, Roeters, and van der Lippe 2016). This makes 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 a general child expenditure function that encapsulates both greedy work and greedy kids. The only assumption we make is that child expenditure increases convexly with hours worked, which can capture a multitude of greedy interpretations: that childcare is more expensive outside of normal working hours, that childcare for those hours is more stressful to organise, that a woman may feel a cost from being away from her children for long periods, or, directly capturing Goldin (2014), that missing out on those crucial hours can result in a wage cost. Crucially, we allow flexibility at work to reduce this expenditure.

We show that, without making any further assumptions, when flexibility increases, the probability of having a child increases for all women. In addition, under the relatively simple assumption of supermodularity between hours worked and flexibility in the child expenditure function, fertility will increase *more* for women who work more hours and hence have higher income. We then micro-found the role of flexibility in a model of time use that introduces multitasking. Flexibility allows a woman to multitask, for example by working at home in the evening when children sleep, effectively raising her total time endowment. This explains the reason an increase in flexibility boosts a woman's fertility: she feels time richer, allowing her to better meet the demands of greedy work and greedy kids. For women who work more and earn more, this fertility effect is stronger.

Our theoretical model makes two key contributions to the literature. First, we place formal structure on Goldin's intuition with a fully fledged career-family model in which women make decisions in an environment where long and particular hours are especially costly in the presence of children. Second, we introduce the novel idea of "greedy kids". We hypothesise that caring for children shares some of the features of a convex wage schedule, by requiring specific hours from the mother.

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. The first lockdown involved severe distancing measures and travel restrictions and led to closure of schools and a number of service industries, with many higher-end occupations moving to online work from home. First, we document a striking and persistent increase in the number of births starting nine months after the first Covid-19 lockdown began in Norway. Using a difference-in-difference monthly event study specification with a cohort x year fixed effects, we show approximately 0.7 additional births per 1000 women per month, or 9.4%, compared to births in the same calendar months three years earlier. Second, we demonstrate the fertility increase to be concentrated among 25-39 year old employed women with a partner.

We use occupational characteristics capturing flexibility from Goldin (2014) to categorise women as having low or high flexibility in the job they held immediately prior to the first lockdown. This categorisation yields a measure of the intensity of the increase in flexibility: low flexibility jobs can be thought of as greedy jobs that had a large flexibility "shock"; 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 those 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. Greedy jobs tend to be associated with higher incomes, and an alternative categorisation based on income confirms that fertility increased most for women earning above median income before Covid-19 lockdown. Alternative specifications, including changing the control year, and introducing placebo lockdowns in other years, confirm the robustness of our empirical findings.

Our empirical findings provide new evidence that flexibility directly impacts fertility. Existing work highlights the importance of work flexibility for women, but does not show impacts on fertility choices. For example, it has been documented that 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 more temporally flexible (Persson and Rossin-Slater 2022), adjust the flexibility-related amenities of their job when returning after maternity (Felfe 2012) and become more productive if given more flexibility (Angelici and Profeta 2020). Women are also disproportionately more likely to work in the education sector relative to men, in order to have the same holidays as their children (Price and Wasserman 2022). New to this literature, we show that increases in work flexibility boost fertility.

Relatedly, flexibility has been shown to reduce the gender wage gap. The classic Goldin and Katz (2011) work compares pharmacists with other professions, arguing that the gender wage gap has shrunk in pharmacy due to a transition to a more linear wage schedule. In a structural framework, allowing women to switch to more flexible jobs helps close the gender wage and hours gap (Hotz, Johansson, and Karimi 2020), while survey evidence shows a correlation between flexible working arrangements and women's wages (Arntz, Yahmed, and Berlingieri 2022).

We focus on Covid-19 lockdown as a source of change in flexibility, but other work documents broader changes due to the Covid-19 pandemic. The United States saw Covid-19 trigger a combination of a baby bust among foreign-born mothers and baby boom among U.S.-born women, particularly young and childless women, and middle-aged women with college degrees (Bailey, Currie, and Schwandt 2022).[1] The baby boom among middle-aged college-educated women is parallel to our findings for Norway. Similar baby booms were observed in Finland (Nisén et al 2022) and Spain (Cozzani et al 2023). Existing research on fertility in Norway has suggested that Norway may have had a positive fertility response to the Covid-19 pandemic due to its better social security system relative to the rest of the world (Sobotka et al 2023, Lappegård et al 2023).

The paper proceeds as follows. Section 2 outlines a theoretical model of flexibility at work and fertility. Section 3 outlines the Norwegian registry data, shows summary statistics and descriptive plots of births over time in Norway, and explains the empirical approach. Section 4 presents the main results, with evidence for the role of changes in flexibility driving fertility increases, and a series of robustness checks. Section 5 concludes.

2 A Model of Greedy Work and Greedy Kids

In this section, we present a theoretical model of the career-family decision, that builds on the intuition in Goldin (2014). Specifically, Goldin posits that certain jobs reward long hours and particular hours, a feature that she terms "greedy work", but are also 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 over work and fertility 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. To simplify the exposition, we focus on the extensive margin of fertility, setting $n \in \{0, 1\}$.

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

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

[1] In the U.S., the baby bust was more severe in states with higher caseloads, while the fertility recovery was stronger in areas with a better labor market recovery (Kearney and Levine 2023).

where z measures the costs associated with children. In the classical economic treatment of fertility (e.g., Becker and Lewis 1973) z is often viewed as being determined by the market prices of child quantity/quality. 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. 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 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. Concretely, mothers who work long hours might enjoy fewer career benefits and promotions than an equally skilled childless woman. Indeed, Goldin 2014 emphasizes that success in greedy jobs is often associated with commitments to work during antisocial hours and to frequent travel, which mothers may be less able to provide. In our empirical setting in Norway, these features are particularly salient because childcare outside of normal working hours is difficult and costly to arrange.[2]

Second, we assume that $\frac{\partial^2 q}{\partial \theta \partial h} \leq 0$, meaning that child-related costs 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 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 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 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.

Lemma 1 *The woman's maximization problem and her indirect utility, conditional on having n*

[2] Childcare during normal working hours in Norway is mostly provided by the state, and is almost free, while the private market for nannies is much less developed than in the US.

[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 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, 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 $V(0,\theta) \geq V(1,\theta)$.

Figure 1 shows the woman's problem in a simple case in which 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 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 utility cost of 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. Because of 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 the illustration 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. 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 (Effects of Flexibility on the Value of Children) *The effect of an increase*

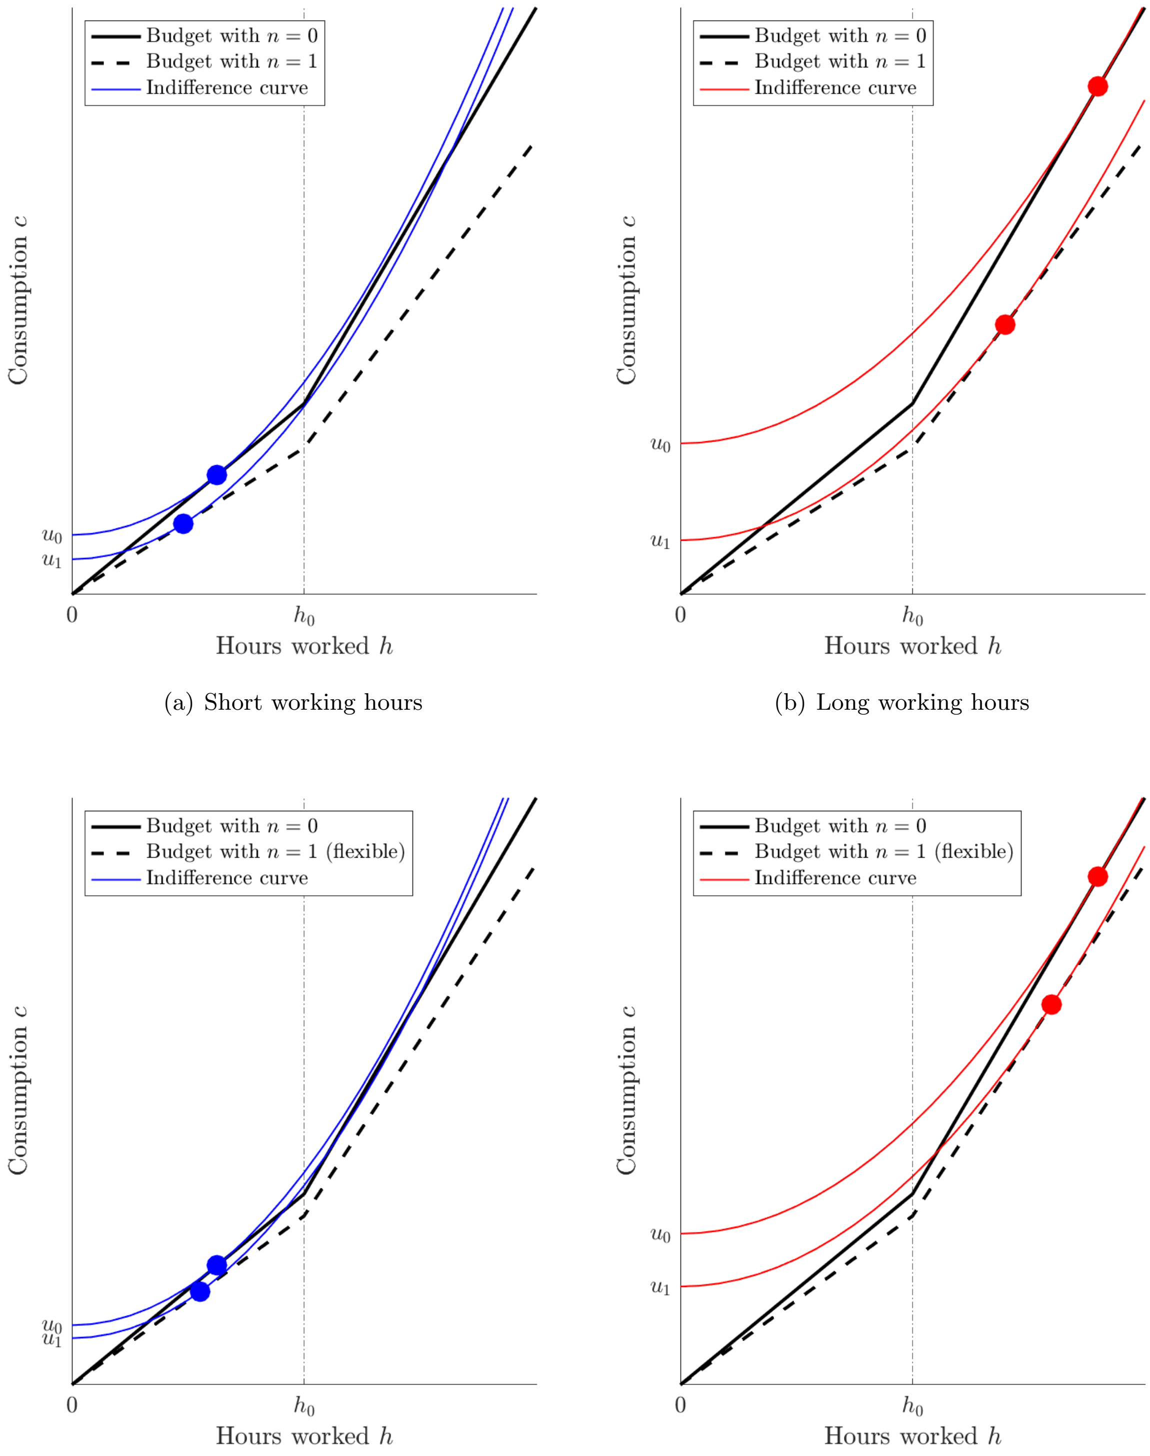

(a) Short working hours

(b) Long working hours

(c) Short hours with flexibility

(d) Long hours with flexibility

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

in the work flexibility parameter θ on the maximized/indirect utility, conditional on having n children, is as follows:

$$\frac{\partial V(n;\theta)}{\partial \theta} = n * \lambda(n) * q'(h^\star(n), \theta) * h^\star(n), \tag{4}$$

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

This result decomposes the marginal benefits of flexibility into four 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 generally 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 an effective cost saving, mothers with high marginal utility of wealth will experience a stronger utility gain. Third, the effect of flexibility is increasing in the marginal child-related costs $q'(.)$. Finally, and importantly for our empirical predictions below, the effect of flexibility is proportional to the number of hours $h^\star(n)$ that the woman optimally chooses to work when she has n children.

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 at least one child, measured by the distance $V(1;\theta) - V(0,\theta)$, increases.*
2. *The incentive to have at least one child increases by a greater amount for women who work longer hours, conditional on the marginal value of wealth $\lambda(n)$.*

The first prediction follows 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.

The second 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.

Our two 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 is 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 panels (b) and (d). This reflects our second empirical prediction, namely, that the incentive to have children increase more strongly for women who work long hours.

We close by remarking that both of our key predictions go through when the wage schedule $y(.)$ and child-related costs $q(.)$ are linear in h. However, strict convexity, as drawn in our illustration, make the predictions more pronounced. On 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 during antisocial 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, to which we turn in the next section. .

3 Empirical Approach

This section describes the data, context and empirical approach for leveraging the first Covid-19 lockdown in Norway as a change in work flexibility for some women, and analysing the 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 (DD) studies described in Section 3.4, we restrict samples to women residing continuously in Norway during the 24-month interval forming the DD-analysis pre- and post-periods of the treatment (lockdown) and control cohorts and who were aged 25-39 at the end of the pre-period. From the employer-employee register, we retained active job records as of March 12, 2017, and March 12,

[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 second prediction would require strict convexity of $q(.)$.

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 Secton 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 of lockdowns and strict regulations on social distancing. The measures led to an immediate reduction in economic activity, and during the first few weeks more than 360,000 people (12% of the labor force) filed for unemployment insurance (UI) benefits (Alstadsæter et al 2020). In most cases, UI benefits covered 65% of lost wages. On March 16, 2020, the Norwegian parliament agreed to temporary changes in the UI program with increased replacement rates, an extension of the maximum duration, and lighter eligibility requirements. 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. Two patterns emerge: first, fertility is declining over this period, consistent with this widespread trend in other high income countries. Second, births are highly seasonal and usually peak in the summer months. Analysed by age, Panel B of Figure 2 illustrates that the declining fertility trend is largely driven by a decline in births to women aged 20-29. In contrast, women aged 30-39 had a relatively stable birth rate during this period.

Even descriptively, 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 seasonal cycle is higher than in the previous year. Fertility stops rising for this age group approximately a year later, coinciding with the end of lockdowns in Norway and a gradual return to work in person.

3.4 Empirical Approach

While the descriptive evidence is convincing of an increase in fertility matching the lockdown period, this could be driven by a number of other trends happening around the same period. To account for these, 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 present in the birth records.

We compare monthly birth probabilities of these women in a 36-month window around Covid-19 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 post-lockdown births would have occurred, and analyse births 12

Figure 2: Births over time

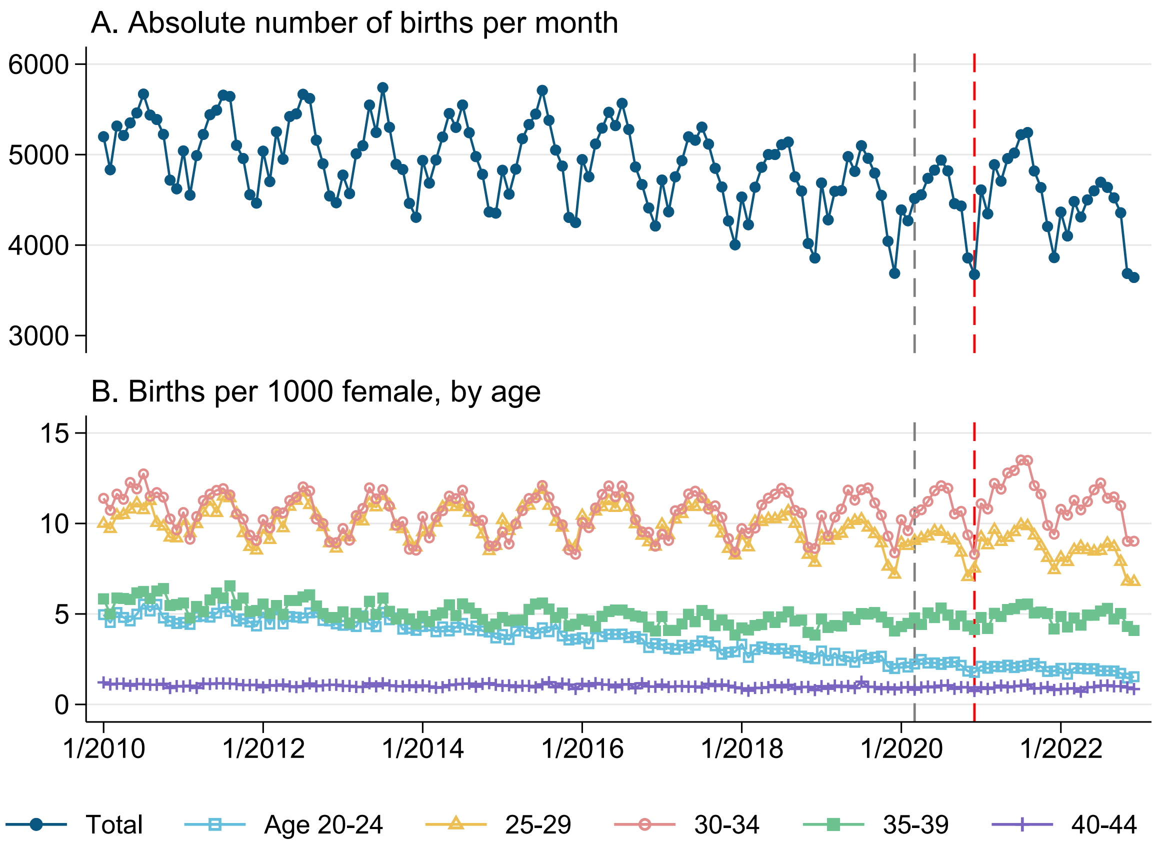

Notes: The grey dashed line indicates the start of the first Covid-19 lockdown, and the red dashed line is placed nine months later, when the first births of lockdown conceptions would have occurred.

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. Any later control year would reduce the number of post months that we could use. Still, as a robustness check, we show in Section 4.3 that changing the control year (with appropriate adjustment of window length to avoid Covid-19 in the control group) does not change our results. 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:

[5]This would assume a woman can conceive again immediately after giving birth, which is a conservative assumption that places an upper bound on the potential births that could happen in any given period.

$$z_{i,t,\tau} = \sum_{\tau=-12}^{\tau=+24} \alpha_\tau Month_{i,\tau} + \sum_{\tau=-12}^{\tau=+24} \beta_\tau PostCovid_i * Month_{i,\tau} + \gamma Cohort_i * Year_{i,t} + \eta_{i,t}, \tag{5}$$

where $z_{i,g,t}$ is the outcome for individual i (birth of a child), 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. $PostCovid_i$ indicates whether the observations are in 2020 (post-Covid) or 2017 (pre-Covid). 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 ($Cohort_i * Year_{i,t}$). 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 lockdown, and after allowing for time-varying cohort effects through cohort * year fixed effects.[6]

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 age group. 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 Covid lockdown and check that birth trends evolved in a similar way in 2017 and 2020, prior to December of those years. 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 birth probabilities are significantly different before and after lockdown, we also estimate a difference-in-difference version of Equation (5) where monthly dummies are replaced with pre- and post-December dummies. 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 A.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 x year fixed effects.

There is a significant and persistent increase in births, ten months after lockdown and onwards. The effect is approximately +0.8 additional births per 1000, relative to a baseline mean of 7.3, so more than 10% of baseline. The effect persists for twelve months from the first lockdown month, with significantly higher births until December 2021, and therefore significantly higher conceptions until March 2021. In total, we estimate around 4,300 additional births during this period. That

[6] We omit individual fixed effects because they are computationally demanding to estimate, and Appendix Figure ?? shows that including individual fixed effects does not change the main estimates.

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. After December 2021, the fertility rate returns to this original trend, parallel with those same months in 2018 and 2019. This is likely because of the end of lockdown policies in March 2021 and a gradual return to business as usual.

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

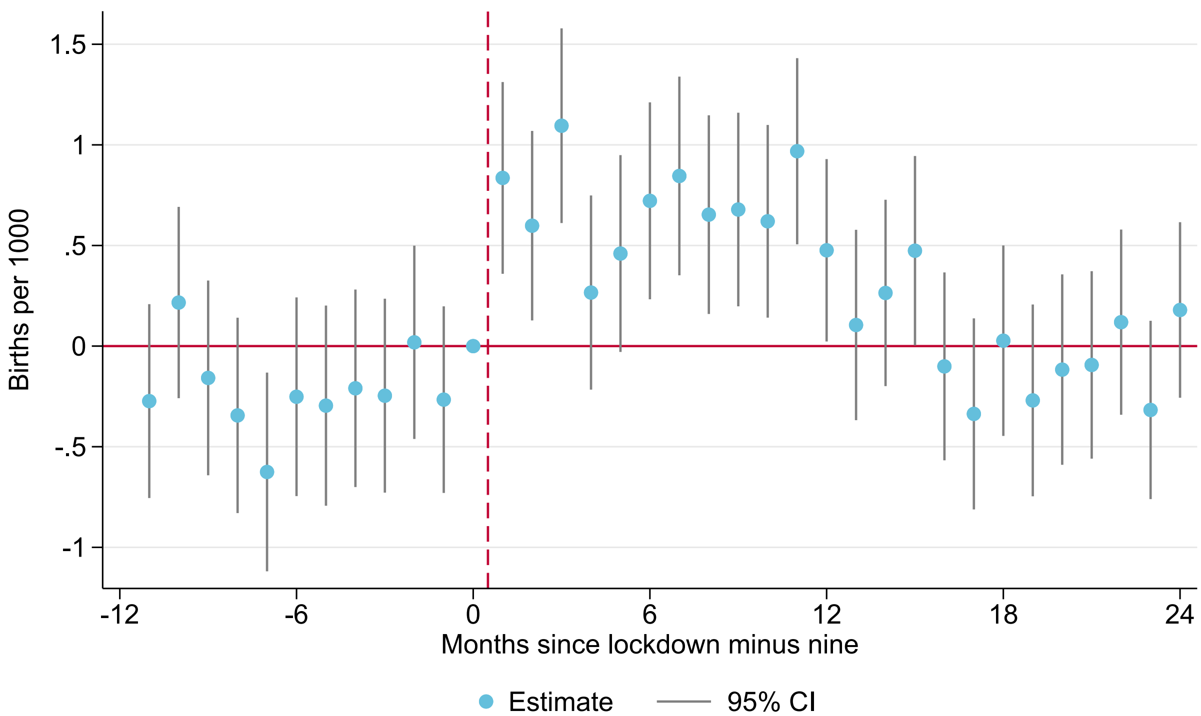

Notes: Treatment cohort consists of women present in Norway Jan 1, 2020, and Dec 31, 2021, and who were 25 to 39 years of age Dec 31, 2020; control cohort analogously defined for women present in Norway Jan 1, 2017. See equation (5) for estimation model.

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 to individuals' homes. We argue that this is the key driver that stimulated aggregate fertility during this period. We show this by demonstrating that the largest fertility responses were evident among women most affected by the movement of jobs into the home. Women most affected by this change are likely to have been employed, with an existing partner, earning more and experiencing a larger increase in flexibility.

First, we show that fertility effects were concentrated among working women with an existing partner. Specifically, in Panels A and B of Figure 4, we compare the fertility response of women employed and unemployed as of March 12, 2020 (relative to the control group of employed or unemployed women March 12, 2017). We find that the aggregate fertility response was entirely driven by women with a job, and that once we focus on these women, the effect on births was larger at +1 additional births per 1000, or 12 percent of baseline. Moreover, for the unemployed group, coefficient estimates reveal imbalance in pregnancies at baseline, with significantly fewer births in April and May of 2020 compared to 2017, 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 the first Covid-19 lockdown.

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

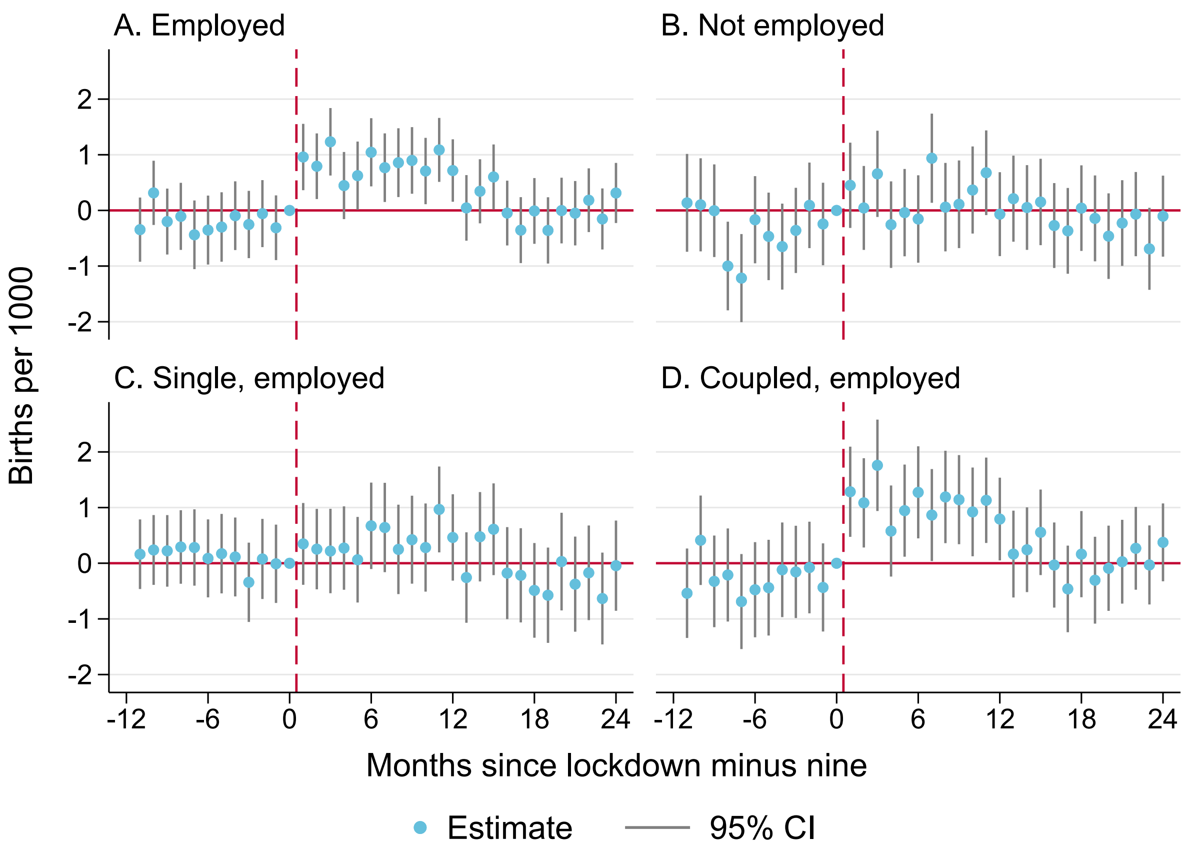

Notes: Employment status measured March 12, 2017 (Control) or 2020 (Treatment), and marital status January 1 of the same year. Being in a couple includes being legally married as well as cohabiting with a partner.

Next, we demonstrate that among coupled, working women, those with the largest increase in flexibility had the largest fertility response. We follow Goldin (2014) and construct a measure of work flexibility using the data collected by O*Net. Goldin 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 split the sample along the median according to occupational flexibility at the time of lockdown.

Low flexibility women experienced a larger flexibility shock, in absolute terms: they were able to keep their career-oriented jobs, but these jobs moved to flexible forms. In contrast, women who already had higher flexibility at work experienced a smaller relative increase in flexibility with lockdown. This intensity of treatment approach, where baseline levels of a variable are used to predict absolute changes and therefore the intensity of a treatment, has been used successfully in Acemoglu and Johnson (2007), Bhalotra, Venkataramani, and Walther (2023) and Ager, Hansen, and Jensen (2017), for example.

The results in Panel B of Figure 5 suggest that fertility increased most for women with relatively low levels of flexibility prior to lockdown, who experienced the larger increase in flexibility when lockdown started. We do not need to rule out positive responses among women with existing higher levels of flexibility (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). In a simple pre-post difference-in-difference comparison of birth outcomes 12 months before and 12 months after December, low flexibility women had a significantly larger birth response than high flexibility women (Appendix Figure A.1).

An alternative categorisation based on income confirms these findings. 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. Among women with below median earnings rank, only three out of the first 12 post-lockdown coefficients are statistically significant (Panel A), while women with above median earnings rank had a statistically significant higher birth probability in eight out of the first 12 post-lockdown months.

Together, these results are empirically consistent with the predictions of Proposition **??** - namely that when flexibility increases, women with higher ex ante earnings, or women experiencing a larger flexibility shock, should have a larger increase in fertility. In the next section, we explore whether other mechanisms could explain our findings 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 earnings and flexibility of work

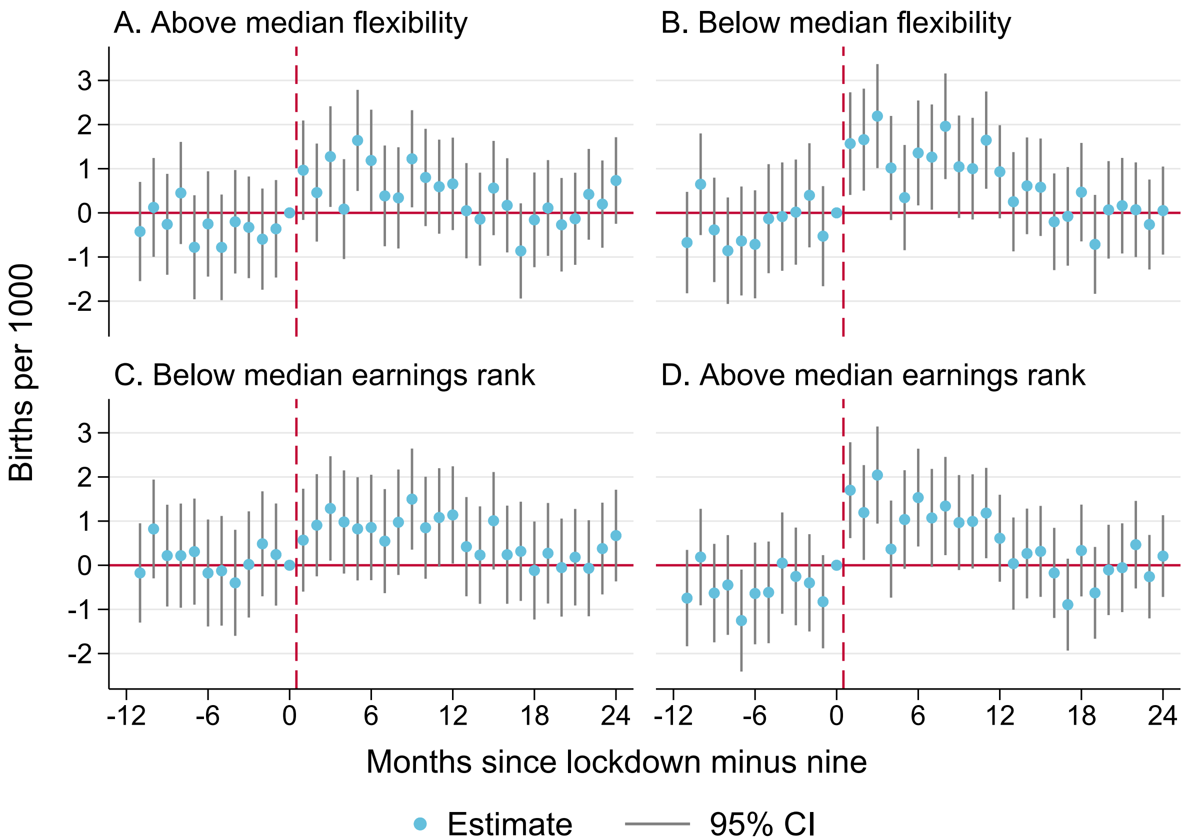

Notes: Occupational flexibility measured as of lockdown using the Goldin (2014) work flexibility scale. Earnings rank, within birth cohort, based on best three of prior ten earnings years (see text).

4.3 Other mechanisms and robustness checks

In our main results, we observe large increases in fertility among high earning, coupled women in Norway, working in relatively inflexible jobs. In this section, we show that this was not driven by changes in income or by the choice of control year, and that a "placebo" lockdown in other years shows null fertility effects.

Theoretically, changes in income or job status are unlikely to explain our findings. 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, Sobotka, Skirbekk, and Philipov 2011), declining during uncertain times and recessions (Schaller, Fishback, and Marquardt 2020). The Norwegian economy shrank by around 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 women who had no lost pay, and those that had at least one month of lost pay.

Our findings are robust to choice of control year. In Panels A and B of Appendix Figure A.2, we show the estimated coefficients from Equation (5) using 2016 or 2018 as control years, respectively. We restrict the post-treatment window to twelve months in each case for comparability. Our main findings are unchanged: in both cases, effects hover around +1 birth per 1000, similar to the effect size for employed, coupled women in Figure 4.

Next, we introduce a "placebo lockdown" in 2016 and 2018. This specification keeps 2017 as the control year, but uses 2016 or 2018 as treatment years with a lockdown assumed to happen in December of those years. Reassuringly, Panels C and D of Appendix Figure A.2 show null effects in both instances. This validates our findings, confirming our results are not driven by other trends or occurrences during this period.

5 Conclusion

While many of the structural drivers of the historical gender wage gap, including education and experience gaps, have disappeared, women continue to earn less than men for the same work. Recent research has laid the blame with flexibility - or lack of it. Specifically, high-paying careers tend to make temporal demands consisting of long and particular hours. These hours 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 have 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 already 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 on average among employed, coupled women aged 25-39. Importantly, we uncover a much larger response among women with less flexible jobs prior to lockdown, and above median earnings.

Our findings do not preclude changes in the magnitude of 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, then 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.

Alstadsæter, A., B. Bratsberg, G. Eielsen, W. Kopczuk, S. Markussen, O. Raaum, and K. Røed (2020). The first weeks of the coronavirus crisis: Who got hit, when and why? Evidence from Norway. *NBER Working Paper* (27131).

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 (2022). The covid-19 baby bump: The unexpected increase in u.s. fertility rates in response to the pandemic. *NBER Working Paper* (30569).

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.

Felfe, C. (2012). The motherhood wage gap: What about job amenities? *Labour Economics 19*(1), 59–67.

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*.

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.

Kearney, M. S. and P. B. Levine (2023). The US Covid-19 Baby bust and rebound. *Journal of Population Economics 36*, 2145–2168.

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.

Lappegå rd, 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.

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 (2022). When dad can stay home: Fathers' workplace flexibility and maternal health. *NBER Working Paper* (25902).

Price, B. M. and M. Wasserman (2022). The summer drop in female employment. *CESifo Working Paper* (9783).

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).

Verhoef, M., A. Roeters, and T. van der Lippe (2016). Couples' work schedules and child-care use in the netherlands. *Journal of Child and Family Studies 25*.

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.

A Additional tables and figures

Figure A.1: The effect of Covid-19 lockdown on births in Norway, a comparison of average birth outcomes in the 12 months before and 12 months after December 2017 and 2020

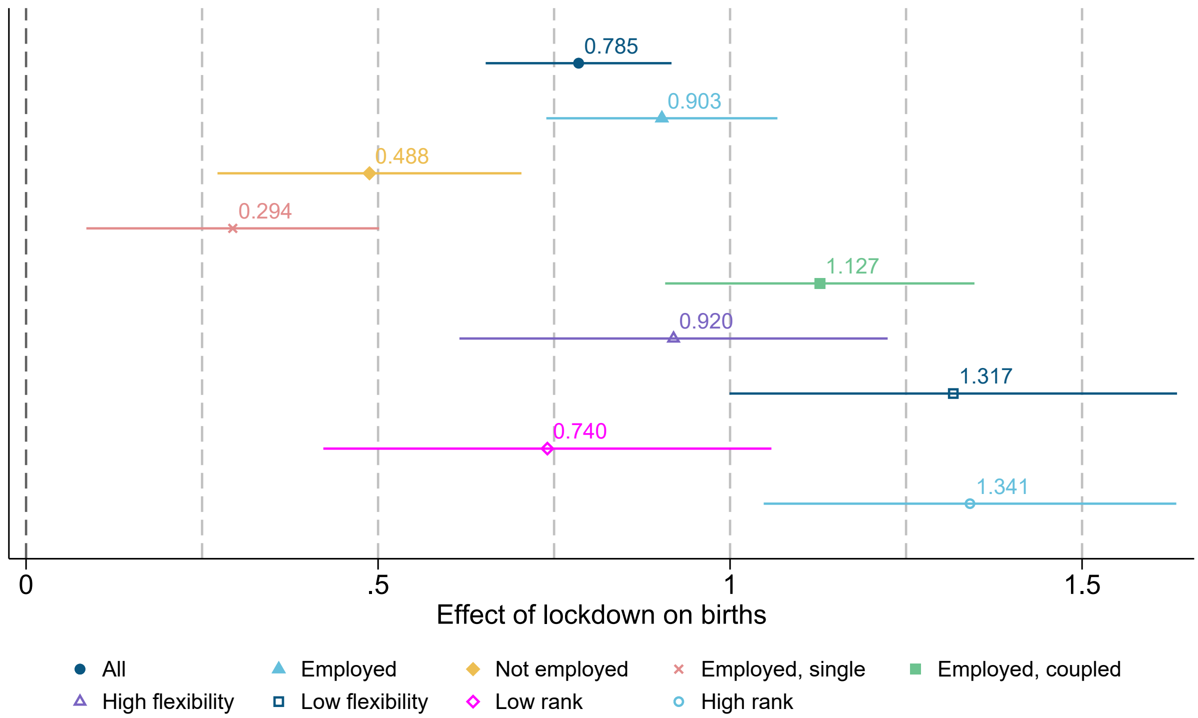

Notes: This figure shows the coefficient on $PostCovid$ in a regression akin to Equation (5), where monthly event study dummies are replaced with pre- and post-December dummies only. The coefficients therefore show the difference in average effects in the 12 months before and 12 months after December 2020, compared to the same differences around December 2017, conditioning on the same control variables and clustering as in Equation (5). P-values from tests of equality of coefficients are: Employed vs. not employed: 0.003; employed, single vs. employed, coupled: 0.000; high flexibility vs. low flexibility: 0.077; low rank vs. high rank: 0.007.

Figure A.2: Alternative control years and Placebo Lockdown

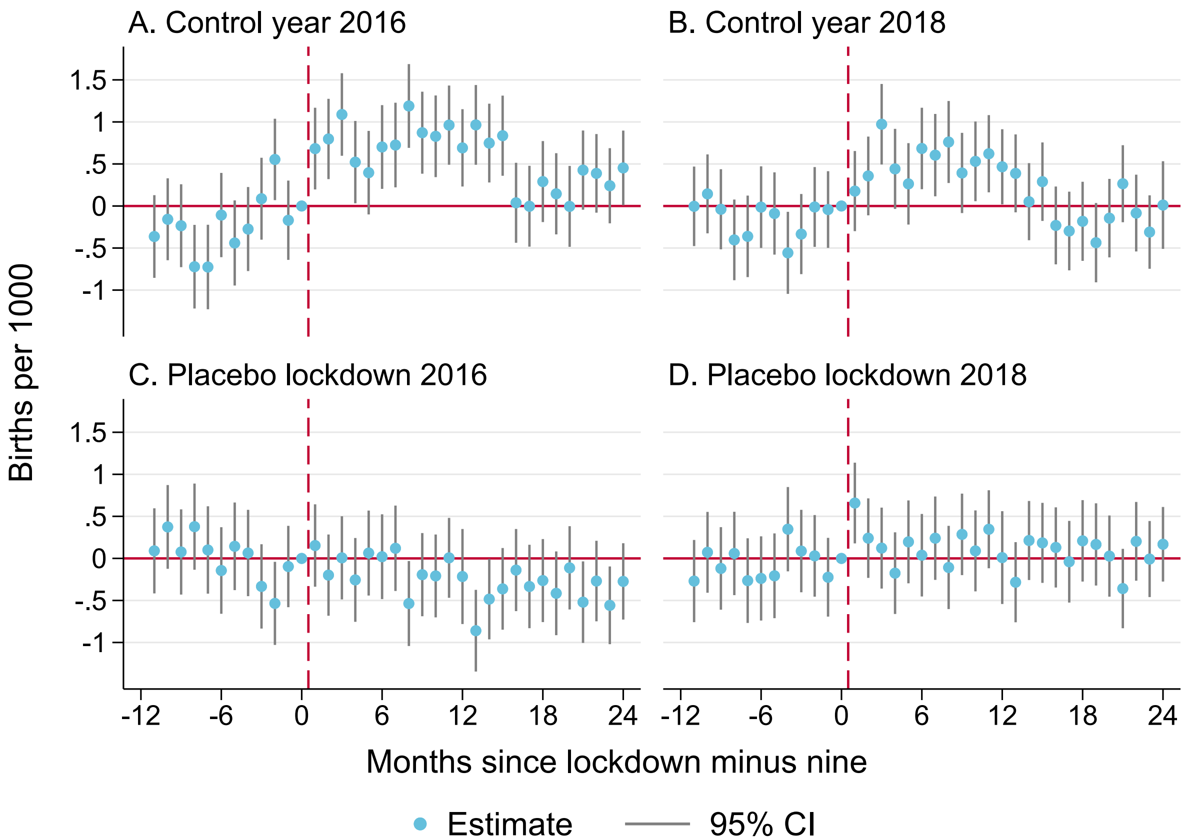

Notes: Panels A and B show alternative control years using 2016 and 2018. Panels C and D show results using placebo lockdowns using only 12 post months