Male Fertility: Facts, Distribution and Drivers of Inequality

Bernt Bratsberg[*] Andreas Kotsadam[†] Selma Walther[‡]

October 5, 2022

Abstract

We document new facts on the distribution of male fertility and its relationship with men's labor market outcomes. Using Norwegian registry data, we uncover a "retreat from fertility": the gap in male childlessness between low and high earners has widened by almost 20 percentage points over the last thirty years, resulting in a remarkable compression of the fertility distribution. Using firm bankruptcies, we show that men experiencing negative labor market shocks are persistently less likely to become fathers and be partnered for at least 15 years after the event. We conclude by documenting that men's fertility penalty to job loss has increased markedly over the last three decades. A challenging labor market fails to shield low income workers with serious implications for family formation.

JEL Classification Codes: J12, J13

Keywords: Male fertility, Unemployment, Inequality

---

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

[†] Frisch Centre. Email: andreas.kotsadam@frisch.uio.no

[‡] Department of Economics, University of Sussex, IZA and CERGE-EI. Email: s.walther@sussex.ac.uk

The authors thank Toke Aidt, Alison Andrew, David Autor, Jesper Bagger, Thomas Baudin, Amalavoyal Chari, Monica Costa Dias, Elisabetta De Cao, Eric French, Paola Giuliano, Paula Gobbi, Chris Hansman, Hamish Low, Andrew Shephard, Vegard Skirbekk, and seminar participants at Oxford, Cambridge, Maastricht, Sussex, CERGE-EI, Norwegian Institute of Public Health, the CReAM (UCL) 20th Anniversary Workshop, the SEHO 2021 Meeting in Boston, the 2022 Alp-Pop conference, the 2022 Leuven Summer Event in Family and Labor Economics, the 1st International Workshop on Family and Migration Economics in Paris, and the Gender and Family Webinar for useful discussion. This work was supported by the Research Council of Norway through its funding of "Intergenerational mobility and labor market inclusion" (project # 300917), "Sustaining the welfare and working life model in a diversified society" (project # 270772) and DIMJOB (project # 296297). Data on loan from Statistics Norway have been essential for the research.

1 Introduction

The livelihoods of low income men have changed dramatically over the past few decades, for the worse. Along with stagnating education levels, they also face falling real wages and employment levels, and growing mortality rates (Case and Deaton 2020). Recent work has documented a "retreat from marriage" in the United States, particularly for low earners (Autor and Wasserman 2013). This stands in contrast to the rising education, falling gender wage gap and increasing economic independence of all women, but also low-income women - changes worthy of celebration (Lundberg, Pollak, and Stearns 2016, Goldin 2021). Among low income households, marginal changes in income can have substantial impacts on children's outcomes, making it crucial to document any changes in the distribution of new births (Løken, Mogstad, and Wiswall 2012).

Using Norwegian registry data, which provides data on all births to the entire population since 1967, we show that there has been a "retreat from fertility" among low-income men. The developments in low-income men's labor and marital outcomes could encompass both changes in out-of-wedlock fertility (low-income men no longer reside with their children) and changes in total fertility (low-income men are no longer having children). We are the first to show that this latter channel is the qualitatively and quantitatively important one, and to analyze its economic origins.

Two new stylized facts emerge. First, comparing men by within-cohort earnings rank, we document childlessness rates that rise exponentially as earnings rank falls. In the most recent 1978-80 cohorts, childlessness rates at age 40 are 72% among the bottom 5% of earners, but only 11% among the top 5% earning men. Second, we document that this inequality in fertility has widened over time: the ratio of childlessness between top and bottom earners has grown by 55% over the last 31 cohort years. We show a remarkable compression of the fertility distribution, with fewer men experiencing a larger share of the population's new births. There has been no concomitant rise in fertility delay or out-of-wedlock births.

Among low income women, we find an increase in parenting partners selected from other cohorts and non-Norwegians, supporting the theory of "serial monogamy" as an evolutionary marriage market response to a rising number of wealthy men in the population (de la Croix and Mariani 2016), as well as a rise in mating by income group that offsets some of the reduced availability of "marriageable men". Comparing women's overall fertility patterns to those of Norwegian men, we find the shape of women's fertility to be more consistent with the previous literature and the standard career-family tradeoff (Baudin, de la Croix, and Gobbi 2015, Kleven, Landais, and Soegaard 2019, Adda, Dustmann, and Stevens 2017, Bhalotra, Venkataramani, and Walther 2022). In the remainder of the paper, we therefore focus on understanding the novel economics of male fertility.

In economic terms, the decline in male fertility that we document could be driven either by a weakening in the preference for children among low-income men over the last three decades or, perhaps more plausibly, by changes in the environment. We propose that underlying developments in the labor market, and particularly the stagnating wages of low-income men, are a key driver of these changes. We proceed in two steps. First, we illustrate high-level patterns consistent with this narrative. Even in Norway, a country with a strong social safety net, growth in the real earnings of low-income men has been feeble over the last 31 cohort years, even as the ratio of top to low incomes grew from 5 to 14 over this period. At the same time, low-income men have become increasingly likely to be single at age 40 and form the lowest number of partnerships in the population, in line with the "retreat from marriage" documented for the United States. We uncover a dominant pattern in the data: differences in the extensive margin of family formation across the earnings distribution, as captured by first births and first families, drive aggregate inequality in family life.

In the second step of our argument, we further test this hypothesis by attempting to causally identify the link between men's labor market and fertility outcomes. Using firm bankruptcies as exogenous variation in male employment and earnings, we take an event study approach that conditions on individual and cohort*year fixed effects and includes same-sex siblings as a comparison group (aspects of this empirical approach overlap with Bratsberg, Raaum, and Røed 2018 and Salvanes, Willage, and Willen 2021).[1] We focus on prime-age men over a 23-year period, following their family and labor outcomes from seven years before to 15 years after the bankruptcy event. The causal relationship between earnings and male fertility is challenging to identify, with other potential confounding factors such as health, incarceration or missing data. Although there is little descriptive evidence to support these other channels (for example, only 0.7% of birth records in our sample have "missing fathers"), the firm bankruptcy approach gives us an additional lever to circumvent some of these issues.[2]

The causal estimates echo the descriptive patterns on fertility outcomes: having verified that bankruptcies lead to a significant drop in earnings and employment, exhibiting the classic "dip, drop and recovery" pattern first shown in Jacobson, LaLonde, and Sullivan (1993), we then show substantial persistence in the negative impacts of job loss on family life. Men who experience a bankruptcy have 2.4% fewer children overall, are 4.4% more likely to be childless, and are 6.6% less likely to be partnered. Interestingly, the probability of fathering a child is reduced in the initial six years but does eventually recover, indicating that the negative long run impacts on total fertility stem from "missed births" during the first six years after the bankruptcy that are not compensated for in later life. A back-of-the-envelope calculation indicates that between 43%-48% of the descriptive relationship at age 40 of childlessness and total fertility can be explained by a causal earnings-fertility link.

With the labor market playing a key role in men's family outcomes, a natural explanation for growing inequality in fertility is changes in the labor market over time: we describe this as a growing "fertility penalty". Specifically, we show several pieces of evidence that negative labor market events carry larger penalties in recent years in terms of earnings, and also on family formation. First, we split the data into an early and late period and show that the negative effects of a firm bankruptcy on its workers are more pronounced in later years, with larger reductions in both income and

[1] Also in a similar approach, Rege, Telle, and Votruba (2007) use plant closures in Norway between 1995 and 2000 and find that marriages decreased as a result.

[2] Specifically, this fraction of missing fathers refers to native mothers in the birth cohorts 1950-80 still present in Norway at age 40.

fertility. To improve on statistical power, our second approach estimates a Mincer-style correlation between job loss in the previous year and having a child the following year, conditioning on a wide set of covariates, for each calendar year between 1990-2020. We show a clear negative trend in this relationship: while men losing their job are less likely to experience the birth of a child in the following year, the crucial finding is that the magnitude of this negative effect has become larger over time. Shedding light on mechanisms, we show that over time, job loss has led to both larger earnings losses and weaker re-employment prospects. Importantly, we verify that these patterns are not explained by changes in the composition of the unemployed group over time. The growing inequality in fertility and family formation can be explained by a more challenging labor market that fails to shield vulnerable workers, with fundamental repercussions for family formation.

We confirm that our main findings on the impact of bankruptcies are robust to a number of different checks: for example, we estimate a specification with family*year fixed effects, which allows for differential trends over time in outcomes across sets of siblings, with unchanged results. We also show that our estimates are similar after taking account of the recent concern over heterogeneous treatment effects in combination with including already treated observations (see, e.g., Goodman-Bacon 2021, Callaway and SantAnna 2021, and Sun and Abraham 2021), with a stacked regression design producing similar coefficients. We also discuss alternative samples, investigate pre-event trends in outcomes in different samples, alternative definitions of firm closures and the removal of bankruptcies that may have occurred outside our sampling window. Our conclusions are robust to these checks.

We contribute most closely to the emerging literature on the economic and family outcomes of low-income men. Case and Deaton's (2020) landmark work describes the loss of a "way of life" among the working class, encompassing changes in the labor market, community, family life and health. A series of papers explores the impact of men's earnings on marital and fertility outcomes in the United States, identifying impacts through data on women's fertility or local birth rates. In a seminal paper, Autor, Dorn, and Hanson (2019) use Chinese import shocks to show that reductions in males' relative earnings increase single motherhood and male premature mortality, and reduce male marriage and fertility.[3] Also focusing on the relationship between male income and fertility, Kearney and Wilson (2018) explore the impact of male earnings growth on fertility and marriage using fracking booms in the U.S., and find an increase in both marital and non-marital childbearing. Finally, Black, Kolesnikova, Sanders, and Taylor (2013) focus on county level birth rates and census data on women's childbearing to estimate that an increase in men's earnings due to an exogenous shock in region-specific demand for coal led to more births. The use of data on women's fertility or births in this series of papers limits their ability to comment on the distribution of births, a

[3] Related to Autor, Dorn, and Hanson (2019), there are several studies that use a shift-share approach to look at family outcomes. Giuntella, Rotunno, and Stella (2021) investigate the effects of trade shocks on marital status and fertility using a household survey in Germany. They find that low educated men working in sectors most affected by increased imports had lower fertility but that marriage rates were unaffected. Similarly, Schaller (2016) and Shenhav (2021) find that lower male earnings reduce fertility and marriage rates. Anelli, Giuntella, and Stella (2019) also use a shift-share approach based on robots to provide evidence that in areas more intensely exposed to robots in the US, new marriages declined, marital fertility declined and out-of-wedlock births increased.

limitation that we are able to overcome.

We draw out two important implications of our findings. First, the shift in the distribution of new births in the population will affect the inter-generational transmission of skills, and children's outcomes in the next generation. In aggregate terms, the average child is more likely to be born to a higher income father among recent cohorts - a positive effect on outcomes via improvements in marginal household income (Løken, Mogstad, and Wiswall 2012). However, the inequality in family stability across men's earnings rank will exacerbate any negative impacts of low income on those children, and particularly boys, who do end up being born to a low income father (Kalil, Mogstad, Rege, and Votruba 2016, Fagereng, Mogstad, and Rønning 2021). Second, there are likely to be welfare implications for low income men who are not participating in family life. Family life is known to socialize men and reduce crime; for example, recent work shows that partner's pregnancy reduces net arrests among men (Massenkoff and Rose 2022). Thus, the deterioration of family formation prospects for low income men may have detrimental consequences that extend beyond the individual to society (Sampson, Laub, and Wimer 2006, Forrest and Hay 2011, Craig, Diamond, and Piquero 2013).

Relatedly, our findings confirm that children are a normal good (Del Bono, Weber, and Winter-Ebmer 2012, Huttunen and Kellokumpu 2016, Lindo 2010, Lovenheim and Mumford 2013).[4] Much of this literature focuses on job loss but it usually analyzes fertility outcomes of individuals already in couples. Almås, Kotsadam, Moen, and Røed (2020) and Hart (2015) show that male earnings in Norway correlate with the probability of finding a partner. Hence, it is likely that job loss affects partnering and by focusing on intact couples, the identified effects are limited to a select sample. More broadly, we add to the growing evidence that the theoretically predicted negative income-fertility relationship has reversed in recent decades in high income countries (Doepke, Hannusch, Kindermann, and Tertilt 2022, Jones, Schoonbroodt, and Tertilt 2011, Fox, Kluesener, and Myrskyla 2019).

The paper proceeds as follows. Section 2 discusses the data, Section 3 presents the key stylized facts on male fertility and earnings, Section 4 discusses our empirical strategy of bankruptcies and Section 5 presents the empirical results. In Section 6 we show that the fertility penalty to negative labor market events has grown over time. Section 7 concludes with implications for society.

[4]Del Bono, Weber, and Winter-Ebmer (2012) show that the probability of a woman giving birth declines in response to her job loss due to a firm closure in the private sector in Austria, while they find no effect of men's job loss. Huttunen and Kellokumpu (2016) confirm this result in a sample of Finnish couples, where female job loss due to plant closures reduces fertility but male job loss has no impact. Both share our concerns of possible selection into firms that eventually close and choose appropriate comparison groups to address this possible bias. Focusing on the U.S. and specifically the response of women's fertility to her husband's job loss, Lindo (2010) estimates a decline in total fertility but an acceleration of births, using an individual fixed effects model to account for possible unobservable characteristics that may relate to job loss and outcomes.

2 Norwegian Context and Data

2.1 Norwegian Context

Fertility in Norway, and in the other Nordic countries, has been falling since the 1980s (Comolli et al. 2020). The Norwegian welfare state is characterized by a dual earner norm while at the same time having strong financial incentives for parents to stay at home (Ellingsæter 2006). There are no particular policy developments that would suggest a decline in fertility. To the contrary, based on evidence from quasi experimental studies from various settings, Bergsvik, Fauske, and Hart (2020) argue that the policy developments in Norway would have led to increased fertility all else equal. They point to increased access to and reduced price of childcare as well as a generous cash for care policy. There are, however, other changes in society over time. Kitterød and Rønsen (2013) show that women have started working more and that men have increased the time spent on household work and childcare. Hart (2015) further emphasizes that costs of living have increased and that the Norwegian universal childcare allowance, which is given to all parents, has fallen in real terms. These factors may affect fertility negatively.

Unemployment insurance in Norway is fairly generous, paying 62.4 percent of lost wages (with lower and upper bounds). During our study period, to be eligible for UI benefits the individual had to document involuntary loss of employment and earnings exceeding 1.5 G during the prior calendar year or 3 G over the past three years (where G refers to the base amount of the Norwegian social insurance system, NOK 100 853 in 2020, slightly less than EUR 10 000). The time limit of UI spells is 24 months. No significant labor market policies took place during this period that could have been instrumental for the outcomes we study.

Demographers have a tradition of investigating the relationship between education and fertility using administrative data (Lappegård et al. 2011; Lappegård and Rønsen 2013). For instance, Kravdal and Rindfuss (2008) and Jalovaara et al. (2019) document that the education-fertility gradient has become less negative for women, has remained positive for men, and that the least educated men are most likely to be childless. There has also been demographic research on the correlation between employment outcomes and fertility in Norway. Kravdal (2002) finds a negative correlation between unemployment and fertility for men, but not women, and Hart (2015) shows that the correlation between earnings and fertility has become more positive over time for both men and women.[5]

2.2 Data

Our analysis is made possible by the use of high-quality Norwegian register data. The data cover the entire Norwegian population, including all births to Norwegian men and women since 1967, with data on all cohorts since 1950. The data also include family linkages, educational attainment, and annual labor earnings. We also use data from the matched employer-employee register in combination with data on firms and bankruptcies.

[5] Kolk (2019) shows similar results using Swedish data.

We operate with four different data extracts. In the *Population sample*, used for descriptive analyses, we include the entire population and focus on cumulative fertility outcomes, studying variation in fertility both across the earnings distribution and over time. The data allow us to track fertility and earnings in the age interval 17 through 50 for individuals born between 1950 and 1970, and through age 40 for those born 1950-1980. For these cohorts, we can also link individuals to their parents, allowing for studies of later-life fertility across the distribution of economic status during childhood. In the *Event study treatment sample* focusing on bankruptcies, we restrict the sample to individuals working in a private-sector firm two years ahead of the firm filing for bankruptcy between 1995-2015, and who were aged 25-35 at the time of the event. For each individual in the event study sample, we stacked their annual outcomes covering the period spanning seven years before and up to fifteen years after the bankruptcy event. To form the basis for counterfactual analysis, we next extracted from the underlying register data siblings of individuals in the event study sample, using similar sampling criteria for the job but with the important exception that the sibling did not work for a firm with a bankruptcy filing during the observation period. This sample selection is discussed in more detail in Section 4. We label this the *Event study control sample*. For the purpose of balanced analysis, we restrict the event study treatment and control samples to families represented with same-sex siblings in both samples (i.e. brothers).

Finally, in the *Stacked cross-sectional samples* used in Section 6, we pool cross-sectional population data for the period 1990-2020 and study changes across time in the correlation between individual unemployment status and outcomes such as fertility, focusing on the age range 25-35 parallel to the event study sample.

In Table 1 we show mean values for the different samples. Note that the table shows means across the whole sample (i.e. all time periods, including pre and post bankruptcy for the treated sample), while in Section 5 we show plots of mean values of all outcomes disaggregated by year. Cumulative fertility is naturally lower and probability of birth higher in the stacked cross-sectional and event study samples than in the population sample, reflecting differences in age of the samples.

A key variable used in later sections is that of registered unemployment. We collect this measure from the register of the welfare administration, implying that the individual has applied for UI benefits at some point during the year. Because a requirement for UI eligibility is involuntary loss of employment, the measure is a fair proxy for individual job loss even though it fails to capture workers who find a new job without seeking UI benefits between jobs.[6] In our population sample of men, about 7 percent were registered as unemployed in a given year.

3 Stylized Facts

We begin by documenting patterns of fertility across time, and heterogeneity in the population, using data on all Norwegian cohorts since 1950. We focus on individuals, and specifically men, who remained present in the country at age 40 and make use of data on their outcomes from 1967

[6]Bratsberg, Raaum, and Røed (2018) estimate that, among native workers, fully 56.5 percent of those who lose their job find new employment without an interim period of enrollment in the UI system.

onwards. In particular, we are interested in how the probability of being childless and total fertility varies with within-cohort earnings rank, and how these patterns have changed over recent decades. We then explore potential mechanisms by studying how real earnings have changed over time in absolute and relative terms, compare this to women, and provide a detailed analysis of changes in partnering, finishing the section with a discussion of the role of other mechanisms such as health, data quality and incarceration.

Table 1: Descriptive statistics, male samples

	Population (1)	Stacked cross-sectional (2)	Event study (treated siblings) (3)	Event study (control siblings) (4)
Outcomes				
Childless	0.213 [0.410]	0.554 [0.497]	0.434 [0.496]	0.433 [0.496]
Birth	0.041 [0.199]	0.098 [0.297]	0.074 [0.262]	0.075 [0.264]
First birth	0.009 [0.095]	0.046 [0.209]	0.032 [0.176]	0.032 [0.175]
Children	1.749 [1.207]	0.757 [0.984]	1.094 [1.167]	1.119 [1.187]
Single	0.325 [0.468]	0.608 [0.488]	0.552 [0.497]	0.531 [0.499]
Unemployed (during year)	0.070 [0.256]	0.141 [0.348]	0.177 [0.381]	0.128 [0.334]
Lifetime earnings rank	50.5 [28.9]	50.1 [26.0]	45.6 [22.4]	50.7 [24.3]
Other characteristics				
Age	40.0 [.]	30.0 [3.2]	33.0 [7.0]	33.3 [7.1]
Education (years)	13.3 [2.7]	13.3 [2.5]	12.7 [2.0]	13.0 [2.3]
IQ	100.5 [13.4]	100.7 [13.4]	98.5 [13.0]	99.2 [13.4]
BMI	21.9 [2.7]	22.3 [3.1]	22.3 [3.1]	22.3 [3.0]
Father's lifetime earnings rank	50.5 [28.9]	51.8 [23.1]	48.0 [22.1]	48.0 [22.0]
Birth year	1964.6 [8.7]	1974.2 [9.5]	1974.2 [5.9]	1974.0 [7.0]
Observation year	2004.6 [8.7]	2004.3 [9.5]	2007.2 [7.2]	2007.2 [7.5]
Age range	40	25-35	18-50	18-50
Observation period	1990-2020	1990-2020	1990-2020	1990-2020
Observations	868 014	9 162 238	267 292	317 579
Individuals			13 087	16 121

Notes: Samples are restricted to men born in Norway to two Norwegian-born parents and present in the country at the end of the observation year. In column 2, unemployment refers to the prior calendar year. Data in columns 3 and 4 limited to individuals 25-35 in the year of event (i.e., year of bankruptcy for treated siblings, year of sampling for non-treated siblings), with a job record in the November file of the employer-employee register two years prior to the event, and matched so that the family is represented in both treated and non-treated subsamples; the means are then computed across the whole sample period. For the event study samples, the range of observation years and observation counts refer to fertility outcomes, which are available for the full data period. Other outcomes, such as single status and earnings, are missing for certain years in the beginning or end of the period. IQ and BMI are collected from conscription data and are missing for 8.0 and 2.9 percent of the sample. Standard deviations are shown in brackets.

Our measure of lifetime earnings rank draws on annual earnings from work covering the period 1967 to 2019. To bypass the need for deflation, for each individual we first compute the within gender and birth cohort earnings percentile at each adult age. Next, we take the average of these percentiles over the age span 25 to 40 - the crucial family formation years - and recompute the individual's earnings rank from the distribution of average percentiles: we call this the individual's lifetime earnings rank. We then show how fertility outcomes at age 40 vary with this within-cohort lifetime earnings rank.

As a comparison point, we calculate a similar measure of earnings rank for individuals' fathers.[7] Father's earnings rank may be a better measure of an individual's ex ante labor market opportunities as it removes some of the endogeneity associated with own earnings rank. We consider similar patterns for own and father's earnings rank as indicative that own earnings rank is a good measure of lifetime opportunities on the labor market.

3.1 Two Facts on Male Childlessness and Total Fertility

We begin by showing how measures of male fertility correlate with men's lifetime earning rank, which yields two stylized facts.

Fact 1: Male childlessness is higher among men with lower relative earnings rank Panels A and B of Figure 1 depict the average percentage of individuals who are childless at age 40, by relative earnings rank within cohort (panel A) and relative earnings rank of their father (panel B), for three representative cohort groups: an early (1950-52), middle (1964-66) and late (1978-80) cohort group. The pattern is striking: on average across all cohorts, only 10% of men in the top 5% of the own earnings distribution are childless, while this number jumps to above 60% in the bottom 5%. In the late cohort group, these numbers are 11% and 72% respectively. This shows marked inequality in men's access to family life. Another interesting feature is that the relationship is not linear: rates of childlessness increase more sharply below the 30th percentile of the earnings distribution.

When examining the relationship by father's earning rank, the overall rates of childlessness vary less, but still decline with earnings. Comparing these two figures, it is clear that while men's own earnings rank is more predictive of childlessness than father's rank, the negative relationship is similar.

Panels C and D depict the correlation between total fertility and own and relative earnings rank. The relationships are very similar to those for childlessness: total fertility increases with both own and father's relative earnings rank, with the relationship particularly strong for the bottom 30% of the own earnings distribution.

Fact 2: Inequality in male childlessness across the earnings distribution has increased over time Figure 2 presents the same data but in a different way, in order to analyze how the relationship between earnings and fertility has changed over time. Instead of taking three

[7]For fathers, lifetime rank is computed from earnings between ages 30-60. The data allow us to use the full 31 years of age-specific percentiles for fathers born between 1937 and 1959. For the oldest fathers, rank is based on earnings during their fifties (95 percent of fathers are born 1916 or later yielding at least ten age-specific earnings percentiles in the data).

Figure 1: Fertility across the earnings distribution.

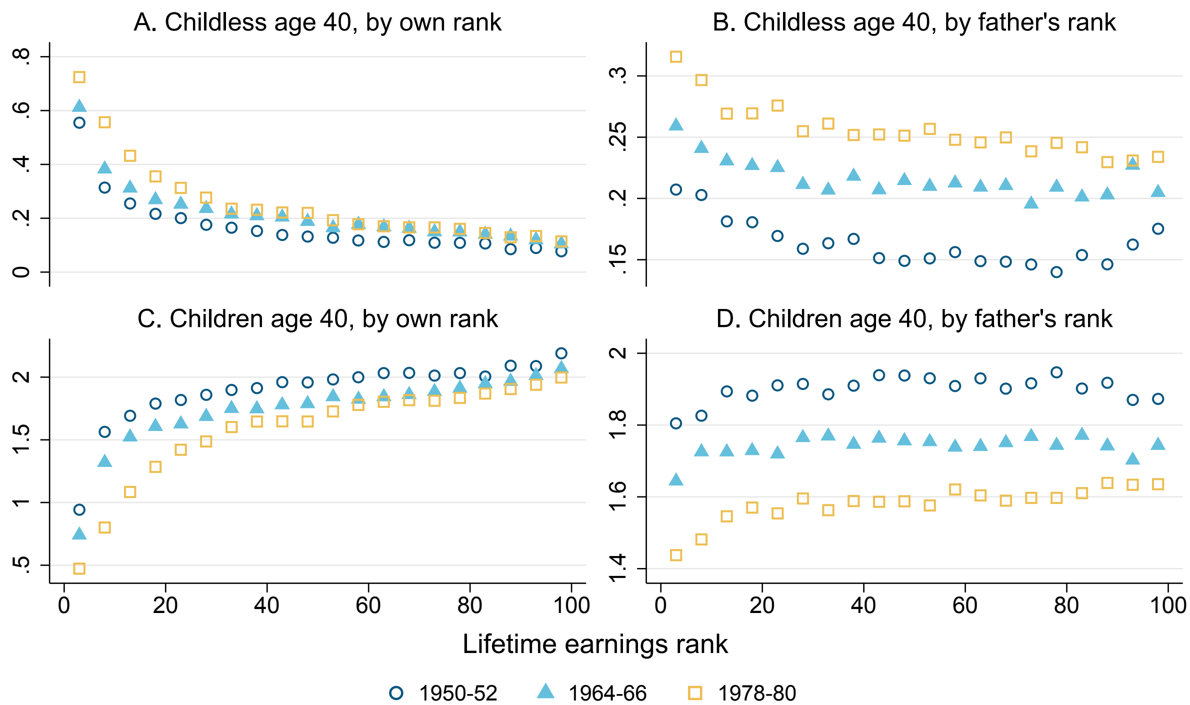

Notes: Each scatter point represents five percent of Norwegian men born between 1950-1952, 1964-1966, and 1978-1980, respectively. Observation count is 244 920.

representative cohort groups, we now take three representative points in the earnings distribution: the bottom, middle and top 10%. We then plot rates of childlessness by cohort, for these three points in the distribution. While childlessness rates have increased for all groups over time, examining childlessness by earnings rank unveils a striking fact: the difference between childlessness rates at the bottom and top of the earnings distribution has widened substantially over time. While the 1950 cohort had a range of 34 percentage points, this soared to 53 percentage points for the 1980 cohort - a marked growth in fertility inequality over this period. This divergence is largely driven by a steep growth in childlessness among the lowest income group, whose trend starts to break from the other two groups at around the 1968 cohort.

Men whose fathers were in the bottom 10% of the earnings distribution have substantially higher rates of childlessness than those whose fathers were in the middle or top of the distribution (Panel B), both of whom have similar, lower rates of childlessness. Panels C and D show these relationships for total fertility. The gap between the total fertility of the lowest and highest earning men has widened over time, from 0.84 children for men born in 1950 to 1.34 children for men born in 1980, largely driven by a steep downward trend in fertility among low income men.

We conduct an additional exercise to check that these trends are not driven by increasing delay in having a child. In Appendix Figure A.5, we show comparable figures with fertility at age 50, rather than age 40. Note that this restricts the number of cohort years we can represent, with 1970 being the youngest cohort present in both the fertility at age 40 and fertility at age 50 figures. Comparing like-for-like cohorts in Figures 2 and A.5, we see that the patterns are very similar, with comparable rates and gaps in childlessness and number of children across the earnings distribution and cohorts. The increase in rates of childlessness among the lowest earners is particularly striking and robust. There is little evidence to suggest that the main patterns are driven by increases in fertility delay among recent cohorts.

Figure 2: Inequality in fertility over time.

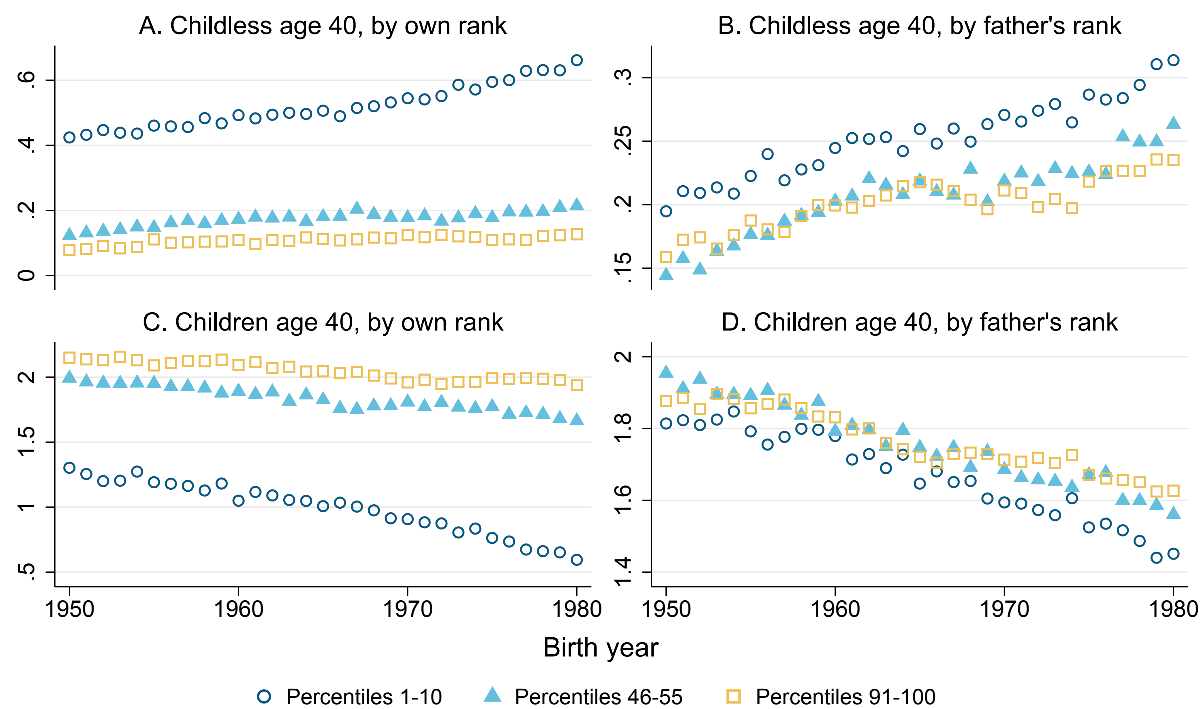

Notes: Scatter points represent ten percent of each cohort of Norwegian men born between 1950 and 1980.

Quantifications The descriptive figures show that the fertility distribution shifted towards the right over the last 31 cohort years, with a greater share of children born to a smaller, higher earnings share of the male population. We can quantify the change in this share. Comparing fertility between the youngest (1980) and oldest (1950) cohorts in our sample, we find that among the oldest cohort, 54158 children were born by age 40 to 86% of men, while in the youngest cohort, 36160 children were born by age 40 to 74% of men. Importantly, this overall decline in fertility masks a striking reallocation of new births from bottom earners to top earners. The bottom 10% of the 1950 cohort had 6.8% of the population's children for that cohort (3691 children), while in 1980, the bottom 10% had only 3.8% of the population's children (1380 children), a reduction of 3 percentage points. The same numbers for the top 10% are 11.2% for the early cohort (6070 children) and 12.4% for the late cohort (4467 children), an increase of 1.2 percentage points. There has been a significant compression of the fertility distribution. This will have implications for the distribution of income and wealth in the next generation (Fagereng, Mogstad, and Rønning 2021).

The extant literature on the retreat from marriage in the US has left the door open for two possible channels vis-à-vis fertility: lower fertility among low-income men, or a growth in the share of low-income men who do not reside with their child. We argue that, in the Norwegian context, the former is the dominant force. To check the extent of out-of-wedlock births across the earnings distribution, we use the fraction of child births where we never see the mother and father on the same household record.[8] According to the records, in the 1980 cohort, 23.6% of births to low earners (bottom decile) were "out of wedlock" compared to 1.6% of births to high earners (top decile). While the fraction of out-of-wedlock births does fall with earnings rank, the overall fractions are low in absolute terms - see Appendix Figure A.6. Certainly, the overall out-of-wedlock rate of 8.0% in the low-income group (23.6% of the 34% of low-income men who start a family) is trumped by the 66% who never have a child. Moreover, the share of out-of-wedlock births has remained relatively stable over time. For example, for the earlier 1970 cohort, 20.4% of births to low-income men were out-of-wedlock, with the same share being 1.7% of high income births. Comparing these out-of-wedlock patterns with those for childlessness rates across cohorts and time, we conclude that a "retreat from fertility" has been the dominant force among low earning men.

3.2 Complementary Facts on Earnings, Partnering and Women

In this section, we present complementary facts on men's earnings inequality, partnering, and the role of women. The descriptive patterns show inequality in male fertility across the earnings distribution, with particular penalties for low income men that have grown over time. A rich literature shows that earnings and labor market potential drive men's value on the marriage market (e.g. Chiappori, Oreffice, and Quintana-Domeque 2012, Bertrand, Kamenica, and Pan 2015), which provides a natural economic link between earnings and male fertility. In the classic Gale-Shapley matching model (Gale and Shapley 1962), a low ranked male (e.g. in terms of earnings) is less likely to find a partner to match with and form a family. We show that this economic mechanism has bite in the data. Specifically, low income men have faced stagnating real earnings even as top earners have seen their real earnings grow sharply, they have the lowest probability of being partnered in the population, and aggregate partnership inequality is driven by the extensive margin - the inequality in having entered family formation across the earnings distribution.

Fact 3a: Earnings of low income men have stagnated over time Real earnings of men in Norway have become more unequal over time. Figure 3 depicts real absolute annual earnings over the age span 25 to 40 in 100.000 NOK by cohort, for the three points of the earnings distribution.

[8] Either they never lived together or they moved in and out during the same calendar year so we miss them in the household files from the end of the year. Data is available for the 1970-1980 cohorts for this measure.

It is clear that while the earnings of men in the top 10% have grown over time, the earnings of men in the bottom 10% have stagnated over time, thus creating widening inequality in income. For the most recent cohort, average earnings for men in the top 10% are 14 times the earnings of men in the bottom 10%, as compared to a multiple of 5 for the earliest cohort in the figure. A similar though less pronounced pattern is seen by father's earnings rank. The marriage market value of men at the lower end of the earnings distribution, as captured by their income, has declined over time in relative terms.

Comparing this to women, the real earnings of women in the bottom 10% of within-cohort earnings rank grew markedly over the same period (Appendix Figure A.7). As a result, the gender wage gap between bottom earning men and women shrunk during this time, from a factor of 10 to a factor of 1.4. Interpreted through the lens of the marriage market, the reservation utility of low income women will have increased, with implications for the partnering prospects of low income men, which we turn to next.

Figure 3: Absolute earnings over time.

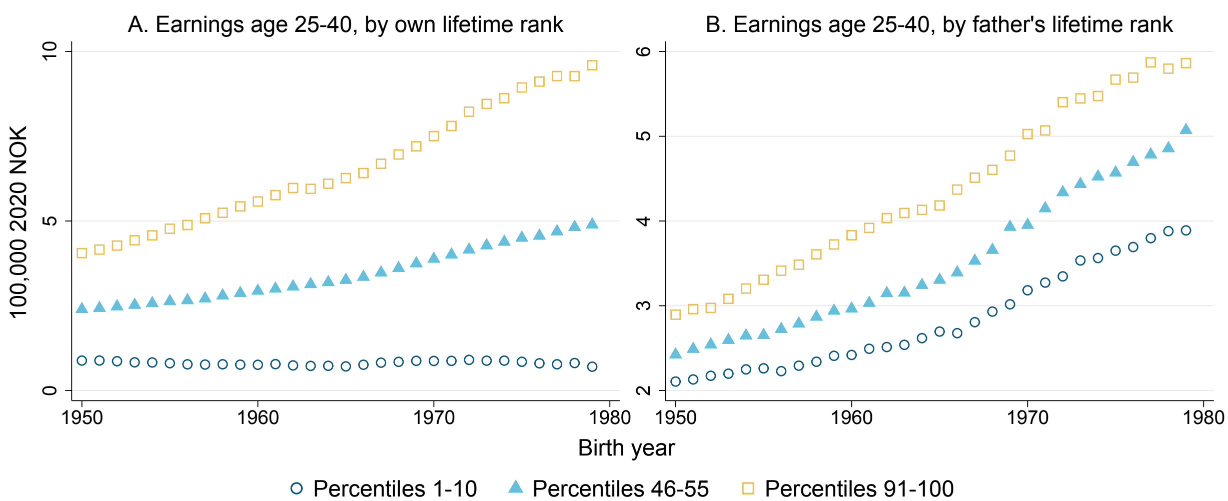

Notes: Scatter points represent ten percent of each cohort of Norwegian men born between 1950 and 1979. Depicted on the vertical axis are average annual earnings inflated to 2020 NOK, in units of 100 000. Observation count is 260 247.

Fact 3b: Low income men are more likely to be single and form fewer families Figure 4 explores how men's partnering relates to earnings rank. In Panel A, we plot the average proportion of men who are single at age 40 by their position in the lifetime earnings distribution, for the three birth cohorts in the beginning, middle and end of our observation period.[9] There are strong

[9]We define single status as one minus the following states: married and living with the marital partner, or cohabiting with someone with whom you have a child. This is for data availability reasons: cohabitations without own children, and cohabitations with others' children, are only recorded from 2004 onwards. To check that the patterns for single status in Figure 4 are not driven by incorrect attributions of childless cohabitations to single status, we compare the rates of single status according to our definition using the old data, and using the richer data

similarities between the patterns seen here and for fertility: single status has increased over time, rates of single status are the highest for those in the bottom of the earnings distribution, and the gap between the top and bottom has widened over time. These patterns are also evident, though the magnitudes are lower, by father's earnings rank in Panel B.

An alternative way of measuring the ability to find a partner is to compare the average number of families formed, across the earnings distribution. We count the number of unique families that resulted in a child that every man in the data has formed by age 40.[10] Panel C of the same figure shows that the average number of partners with whom one has formed a family by age 40 is strongly correlated with earnings. More surprisingly, this gap in the average number of partners between the lowest and highest earning men has widened over time, similar to the widening inequality in fertility by earnings.

What has been the impact of this reduced partnering of low income men on women? Appendix Figure A.9 shows childlessness and total fertility figures for women that are comparable to Figure 1. A caveat applies when interpreting these patterns, which is that lifetime earnings rank will have a different meaning for women compared to men, as women often take time out of work when having children. Further, the length of time women take off work when having children has changed across cohorts. Other factors, such as the career-family tradeoff and the fertility penalty, are of pivotal importance in explaining the earnings-fertility link for women, and have been well-studied in the literature (e.g. Kleven, Landais, and Soegaard 2019, Adda, Dustmann, and Stevens 2017, Bhalotra, Venkataramani, and Walther 2022).

Panel A in Figure A.9 crystallizes this distinction in the earnings-fertility link between men and women: the relationship between women's childlessness rates and relative earnings rank is U-shaped, rather than exponentially declining: female childlessness is highest at the extreme ends of the earnings distribution. This confirms the findings of Baudin, de la Croix, and Gobbi (2015) for the United States, who show that childlessness rates are highest for women with lowest education and highest education levels, arguing a social poverty mechanism for the lowest and an opportunity cost mechanism for the highest.

Women's value on the marriage market is less likely to be driven by earnings, and more likely to be driven by factors associated with fecundity, such as age (see, e.g., Low 2022, Andrew and Adams-Prassl 2022). Consistent with non-income drivers of women's marriage market value, we note that rates of childlessness do not vary across the earnings distribution for women nearly as

from 2004 onwards, across the points of the earnings distribution. Figure A.8 in the Appendix summarizes the data. A few important points emerge. First, the middle and top percentiles of the earnings distribution are essentially flat over cohorts in terms of rates of single status at age 40, regardless of which data source is used. Second, for the low earners, according to our definition, single status has grown from 73.25% to 83.74% across cohorts, while the richer data indicates a growth from 67.05% to 75.08%. Thus, our definition overstates single status among low earners by around 6-8 percentage points, and the growth in single status by around 2 percentage points. We do not think these magnitudes are large enough to dominate the overall patterns seen in Figure 4, though we do acknowledge them to be an upper bound. Third, we check which "omitted" category drives this difference, and find most of the difference to be due to cohabitations without own children, which represented 5.99% of low earners in the most recent cohort.

[10]Multi-partner fertility, using the same definition as ours, is also discussed in Lappegård and Rønsen (2013) for Norway.

Figure 4: Marital status and number of partners.

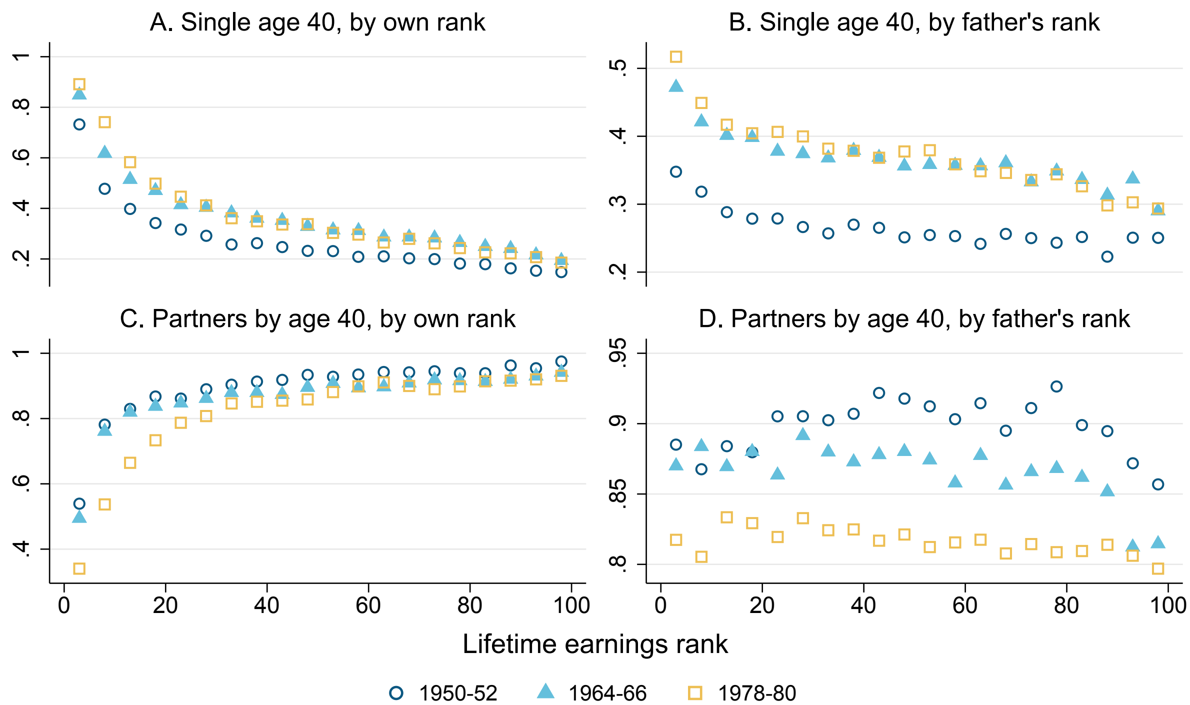

Notes: Each scatter point represents five percent of Norwegian men born between 1950-1952, 1964-1966, and 1978-1980, respectively. Panels C and D count the the number of unique partners with whom the male has fathered a child, including zero for those childless at 40. Observation count is 244 920.

much as they do for men.

While we see a rise in childlessness among low income women across cohorts, likely due to a reduced availability of "marriageable men" in comparable income groups (Autor and Wasserman 2013), low income women have lower rates of childlessness than low income men. The natural question is, who are they having these children with? To answer this question, we expand our dataset to include also non-Norwegians. A simple adding-up calculation reveals that of those women born in 1980 in the bottom 10% of earnings with at least one child, 86% of them have children with a native partner, of which some are across-cohort matches; 11% of these women have a child with a non-Norwegian partner, and the remaining 3% have a child where the father's name is missing on the birth certificate. We also detect a rise in mating by income, which offsets some of the reduction in available low income male partners. Therefore, low income men retreating from fertility are being replaced with a combination of Norwegian men from different cohorts, non-Norwegian men, and a rise in childlessness among low income women.

Fact 3c: The extensive margin of partnership drives aggregate patterns Next, we explore whether the aggregate relationships between men's earnings rank and family formation mask heterogeneity between the intensive and extensive margins of fertility and partnering. That is, the negative relationship between earnings and fertility, and positive relationship between earnings and number of partners, could be driven by the extensive margin of having a first child and finding a partner to have a child with *at all*, or the intensive margin of having more than one child, and more than one partner.

Panel A in Figure 5 shows the *intensive margin of partnership* - the average number of families formed, conditional on forming at least one. We see a striking reversal of the positive relationship between overall partnering and earnings seen in Figure 4. Specifically, conditional on having had at least one partner with a child, low-income men are significantly more likely to have formed multiple families, compared to high income men. Thus, the intensive margin of partnership shows an opposing pattern to the overall relationship between partnering and earnings. This unveils an important fact: there has been a compression of the partnering distribution on the extensive margin across cohorts, that mimics the compression of the fertility distribution. A smaller share of the male population is participating in family formation.[11]

This suggests that the overall increase in inequality of partnering across the earnings distribution is driven by a powerful extensive margin i.e. having at least one serious partnership that results in a child. Appendix Figure A.10 confirms this: while rates of multi-partnering are stable across cohorts and within earnings rank, the extensive margin of partnering (the rate of no-partners) has shifted markedly over time. This is especially pronounced among the lowest earners, who have seen an exponential rise in the probability of having no serious partnerships.[12]

We can quantify partnership differences for the 1970-80 cohorts, comparing the bottom 10% and top 10% of earners (see Appendix Figure A.6).[13] Low-income men are less likely to find a partner to have a child with compared to high income men; in the most recent cohort, only 34% of them succeed. Conditional on succeeding, they show relationship instability, fathering children with on average 1.23 partners, in partnerships that last only 4.9 years on average. [14] High income men are more likely to form a family (87% of them are successful), and conditional on forming a family, have fewer overall partners (1.06) and more stable relationships that last a mean of 11.2 years.

This inequality in partnership stability will have implications for the *intensive margin of fertility*. Figure 5 shows the distribution of total fertility conditional on having at least one child against earnings rank in Panel C (and against father's rank in Panel D, with a more muted relationship). Intensive margin fertility is positively correlated with lifetime earnings rank. As well as higher rates of childlessness, low earning men also have fewer children conditional on having at least one,

[11]Figure A.11 confirms that the extensive and intensive margins of partnering move in opposite directions when comparing percentiles of the earnings distribution across cohorts.

[12]In the same Appendix Figure A.10, we also explore whether men who form multiple families are positively or negatively selected based on father's earnings rank, and find a statistically significant difference for just over half of the cohorts in the figure, where multi-partner men appear to be *negatively* selected on father's earnings rank relative to single-partner men.

[13]We can quantify these multi-partnering differences for all cohorts, but need household records to study durations. Household records are first available in 1991.

[14]Durations are censored at age 40 years.

echoing the patterns in Figure 1 Panel C. Perhaps unsurprisingly given their shorter partnership durations, low income men also have fewer children in those partnerships that they do form.

Figure 5: Intensive margin fertility and partnership across the earnings distribution.

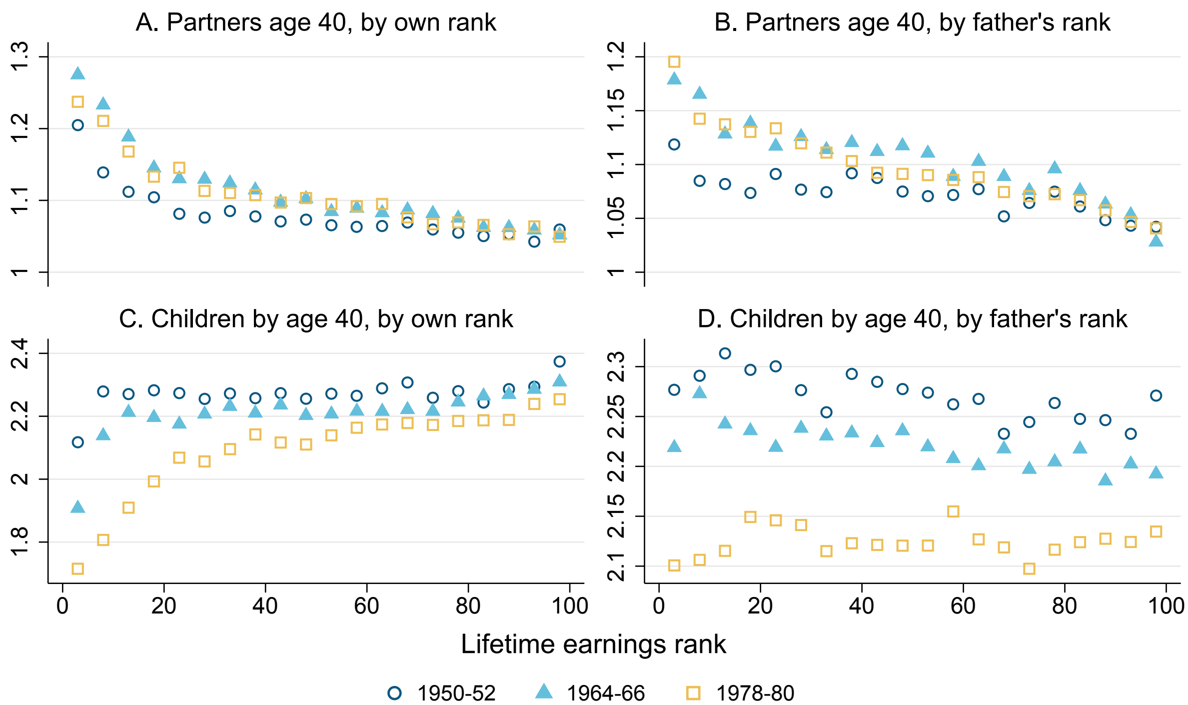

Notes: Population is restricted to those with at least one child by age 40. See also notes to Figures 1 and 4.

Summary and implications of findings We uncover a "retreat from fertility" among low income men without any associated increase in out-of-wedlock births. First, we establish that on the extensive margin, low earning men are more likely to be childless than high earning men, with a gap of 53 percentage points in the most recent cohort. This gap has grown over the last 31 cohort years, with widening inequality in childlessness between low and high earners. The economic link between earnings and fertility via the marriage market is evident in the data: low earnings rank men have faced stagnating real wages and growing inequality in the labor market, as well as reduced partnering and more partnership instability. The widening inequality in family life across the earnings distribution is primarily driven by changes along the extensive margin (low earning men not forming any families).

We draw out two serious implications. First, the change in the distribution of new births across the male population will affect investments in and outcomes of children. Though these changes may be positive in aggregate terms due to the increase in average father's income, the family instability we document among those low income men that do have children will have negative consequences

(Løken, Mogstad, and Wiswall 2012, Kalil, Mogstad, Rege, and Votruba 2016, Fagereng, Mogstad, and Rønning 2021).

Second, there are likely to be welfare implications relating to crime. The literature has shown that family life socializes men and reduces crime, so that the deterioration of men's family formation prospects may have detrimental consequences that extend beyond the individual to society (Sampson, Laub, and Wimer 2006, Forrest and Hay 2011, Craig, Diamond, and Piquero 2013).

Alternative mechanisms In Appendix A.1 we analyze the role of other factors that may correlate with earnings rank and fertility. First, we show correlations of health outcomes such as long-term disability, BMI and height with men's earnings rank, which are modest and do not mimic the rotation we see in the fertilty-earnings relationship across cohorts. Second, we show that while incarceration rates do fall exponentially with earnings rank, the overall fraction of time spent in prison for low earners is too low to explain the high rates of childlessness we document. Third, we focus on data quality, and the possible role of "missing dads". In our dataset of native mothers in the 1950-80 cohorts still present in Norway at age 40, only 0.7% of records do not have a father's name specified. Although we cannot show the relationship between the proportion of birth records with a missing dad and men's relative earnings rank by definition, we show that the same proportion as a function of *mother's* earnings rank is very low overall, and is only 3% for the lowest earning women and less than 1% for the highest earning women. Overall, we do not find reason to believe that any of these alternative factors are a leading explanation for the patterns that we show above.

4 Identifying the Earnings-Male Fertility Link

4.1 Empirical Strategy

Using descriptive data, we document a striking inequality in male fertility across the earnings distribution, and uncover significant heterogeneity in fertility and partnering across the intensive and extensive margins. This indicates an important role for men's earning power in driving inequality in male fertility. Still, these correlations may be confounded by omitted variables that affect both earnings and fertility. In this Section we outline an empirical approach to causally identify the relationship between men's labor market prospects and their fertility. An ideal experiment would "pick up" a man at a random point in the distribution of lifetime earnings rank in Figure 1, and place him a few spaces down the earnings rank, and then follow the impact on his fertility and family formation. To get as close as possible to such an experiment, we use bankruptcies as an exogenous shock to labor market outcomes. We recognize that the analogy to the ideal experiment is not perfect as job loss likely affects other factors that may be important for fertility and partnering such as identity, self-esteem, and time use.

Firm bankruptcies are known to cause increases in unemployment probability and have been used commonly in the literature as a shock to employment prospects, including in Norway (Bratsberg, Raaum, and Røed 2018). In our sample, bankrupt men experience an 8-point fall in their average annual earnings rank, from 53.7 to 45.8.[15] Thus, they have a significant impact on men's earnings.

Bankruptcies are relatively common, with 1% of the Norwegian working population experiencing a bankruptcy in any two years. We verify that bankruptcies definitely lead to business closures: in our sample, by year three after the bankruptcy filing, no individuals previously employed at the to-be-bankrupt firm are still working there. In this sense, and in contrast to using plant closures as in the previous literature estimating the relationship between job loss and fertility (Del Bono, Weber, and Winter-Ebmer 2012, Huttunen and Kellokumpu 2016), we use a measure that in our context is more closely linked to job loss than general firm closures, as we show below. Given that we define treatment as being exposed to a bankruptcy, our estimates can be interpreted as an intention to treat design, where the treatment will entail job and earnings losses, but potentially also changes in self-esteem and identity.

Although bankruptcy filings are associated with a large increase in unemployment risk and reduction in earnings, they may not be purely exogenous because individuals with certain unobservable characteristics may select into financially distressed firms that eventually go bankrupt. If these characteristics also affect their family outcomes, then the estimated impact of bankruptcies on these outcomes may be biased. Alternatively, firms in distress may have slower wage growth than non-distressed firms. This is a well known fact in the layoff literature going back at least to the seminal work of Jacobson, LaLonde, and Sullivan (1993). These selection and compositional concerns are also discussed in Dustmann and Meghir (2005), in particular the choice of when to sample individuals working at a to-be-bankrupt firm, with evidence in favour of $t-1$ and $t-2$ (t being the year of the bankruptcy).

Our approach makes use of within-individual time variation in exposure to the shock. We follow the approach of Jacobson, LaLonde, and Sullivan (1993), who empirically estimate the labor market penalty to job loss following, in their case, mass layoffs. Their main empirical specification is an event study design with individual fixed effects (also used by Lindo 2010) and a comparison group of workers who are not laid off. The individual fixed effects approach removes any bias arising from time-invariant unobservable characteristics that correlate with both exposure to mass layoffs and the set of outcomes, such as IQ and height. Jacobson, LaLonde, and Sullivan (1993) note that even when allowing for specific trends and a comparison group, the impact of mass layoffs is evident as a dip in earnings already three years before the event, but not before then.

We improve on their specification in two ways. First, by using bankruptcies, we circumvent some of the issues of mass layoffs, where firms may displace lower productivity workers, thus creating a select sample among the treated. Bankruptcies result in job loss among all workers employed in the firm at the time. Second, to address the issue of time-varying characteristics that might affect labor market and fertility outcomes, we also include a control group, using matched same-sex siblings working in a firm that does not go bankrupt. The idea is that brothers provide a closer

[15] This is estimated in a difference-in-difference regression; see details in Table 3.

comparison group of the counterfactual path of labor market and fertility outcomes for the treated group. Specifically, a possible confounding situation where men on a declining earnings path, and a declining fertility path, take a job at a distressed firm that eventually goes bankrupt, is better addressed by siblings rather than a random group of control men. Our sample of control brothers are chosen to match the bankruptcy sample as closely as possible, with the same age range and year range, and we draw a random sequence of years from the sample year range.[16]

We follow Dustmann and Meghir (2005) and choose to sample individuals employed at the eventually-bankrupt firm two years prior to the bankruptcy. Choosing an earlier year improves the exogeneity of workers being attached to a particular firm, but reduces the exposure of the individual to the bankruptcy because individuals are more likely to have left the firm by the time the bankruptcy occurs. Therefore, the choice of two years prior provides a balance between these two trade-offs. We also conduct further robustness checks on this assumption in Section 5.5, by changing the timing of when we sample individuals.

The estimating equation is:

$$z_{i,g,t} = \sum_{\tau=-7}^{\tau=+15} \alpha_\tau Time_{i,\tau} + \sum_{\tau=-7}^{\tau=+15} \beta_\tau Treat_{i,g} * Time_{i,\tau} + \theta_i + \gamma Age_{i,t} * Year_{i,t} + \eta_{i,g,t}, \qquad (1)$$

where $z_{i,g,t}$ is the outcome for individual i, and where g denotes firm and t observation year. $Time_{i,\tau}$ is a dummy variable representing time around the event year, and $Treat_{i,g}$ indicates whether the firm g of employment at time -2 goes bankrupt two years later.[17] The coefficient β_τ gives the differential impact as compared to the sibling trajectories captured in $\alpha_\tau Time_{i,t}$. As well as including individual fixed effects, we also include a full set of cohort * year fixed effects. The data is centered so that bankruptcies occur at time zero. Standard errors are clustered at the firm level (i.e., the workforce of the individual's employer at time -2). Identification from the above estimating equation relies on siblings providing a valid counterfactual trajectory for the outcomes of treated individuals, had they not experienced the bankruptcy, and after allowing for individual time-invariant differences through individual fixed effects, and time-varying cohort effects through cohort * year fixed effects. In a robustness check, we also allow for sibling-specific time trends by including family * year fixed effects.

Our extended time period of analysis, seven years before and 15 years after bankruptcy, allows us to closely monitor the evolution of outcomes before the bankruptcy. As we discuss in the

[16]Key to the sampling design is that, in the base year, the treated sibling holds a job in a firm that will go bankrupt while the workplace of the non-treated sibling does not face bankruptcy. Both siblings may, however, work for employers that file for bankruptcy in other years of the time sequence when we follow the individual. Specifically, in our control group sample, 38 men in year -1 and 15 in year 0 work in a firm that goes bankrupt at time 0. In a robustness check in Section 5, we address the concern that bankruptcies in the control group may contaminate the design, and show that dropping these individuals does not change our estimates.

[17]Not all data are available for all outcome years, but 52 percent of the event study sample can be followed for the full 23-year window. In addition, some observations are dropped due to deaths (0.55 percent) and emigration (0.75 percent). Finally, the sample is restricted to men with brothers. We check and confirm that the results are robust to using the subsample that can be followed all years (the balanced sample - see Section 5.5).

results section, and consistent with Dustmann and Meghir (2005), Huttunen, Møen, and Salvanes (2011), and Jacobson, LaLonde, and Sullivan (1993), there is evidence of some effects of bankruptcy starting already a few years prior to bankruptcy filing, but we see no effects at five years or earlier. Therefore, we set the omitted year to be -5, so that all coefficients are estimated relative to the mean outcomes in this initial year. We do this purposefully because our goal is to draw comparison with outcomes unaffected by the treatment. In Section 5.5, we also show that pre-event trends in outcomes across various sampling groups are reassuringly similar, with our chosen control group performing much better than alternative samples in tracking the pre-bankruptcy outcomes of treated individuals. We also consider alternative definitions of firm closures, remove bankruptcies that may have occurred outside our sampling window, consider an alternative control group of men experiencing bankruptcies in the future (following Fadlon and Nielsen 2019 and Salvanes, Willage, and Willen 2021) and conduct a stacked regression design to allow for heterogeneous treatment effects over time. We find that our conclusions are robust to these checks.

We consider impacts on a wide range of time-varying outcomes, including: unemployment status, log earnings, whether an individual experienced the birth of a child, whether the birth was the first child, total (cumulative) fertility, and whether an individual is single (unmarried or unpartnered).[18] We also check the impact on disability status, a common way of claiming benefits in Norway after unemployment benefits expire. As well as providing coefficient plots of the difference-in-difference event studies, we also conduct a simple pre and post difference-in-difference estimation to check that average outcome levels are significantly different post-treatment (see Table 3).

4.2 Descriptive Statistics

The descriptive statistics in Table 1 show that the men we use for the event study design (which draws on a younger segment of the population than that used in our main figures) are somewhat less educated and have fathers of a lower earnings rank compared to the population average in column (1). Comparing the treatment and control samples, we see that they are naturally identical on father's earnings rank. They are also similar, but not identical, on other aspects that are measured pre-treatment such as educational attainment and IQ. This is why it is important to add individual fixed effects to the estimation. In addition, the treated brothers are slightly younger, which we account for with cohort * year fixed effects. Outcomes such as unemployment and fertility are reported in this table as a sample averages across all time periods including pre-bankruptcy and post-bankruptcy, so the overall control sample has significantly higher fertility and better labor market outcomes. Figures 6 and 8 discussed below disaggregate these descriptive average outcomes by year.

In Table 2 we investigate the differences in the characteristics of the firms in the treated and control samples. We see that the males in the treated sample work in smaller and younger firms

[18] We do not use lifetime earnings rank as an outcome because it is not possible to create a clean definition of this that does not encompass both the pre- and post-bankruptcy periods. However, we note that any impact on annual earnings will clearly affect lifetime earnings rank.

than their brothers in the control sample. Digging deeper into the firms that these individuals work for, we see that, reassuringly, the three most common industries in the bankruptcy sample (construction, manufacturing and retail/wholesale trade) coincide with the three most common industries in the non-bankruptcy sample. However, a larger share of the bankruptcy sample works in hotels and restaurants, while public administration and health services are more common in the non-bankruptcy sample.

Table 2: Descriptive statistics, comparing treated and control firms

	Treated (bankrupt) firms	Control (non-bankrupt) firms
	(1)	(2)
Observations	267 292	317 579
Individuals	13 087	16 121
Firms	6 873	8 581
Mean firm size	46.5 [112.1]	1293.3 [3617.3]
Mean firm age	9.6 [8.0]	17.3 [13.3]
Mean firm log wage	5.167 [0.339]	5.353 [0.306]
Manufacturing	0.224 [0.417]	0.195 [0.397]
Construction	0.230 [0.421]	0.165 [0.371]
Retail/wholesale	0.200 [0.400]	0.172 [0.377]
Transportation	0.060 [0.238]	0.086 [0.280]
Hotels/restaurants	0.075 [0.263]	0.023 [0.151]
Info/communications	0.050 [0.217]	0.047 [0.211]
Prof/tech services	0.042 [0.200]	0.039 [0.194]
Admin/support services	0.049 [0.215]	0.048 [0.213]
Public admin	0.000 [.]	0.052 [0.221]
Health services	0.009 [0.097]	0.052 [0.221]
Other	0.062 [0.241]	0.121 [0.326]

Notes: Firm characteristics are measured at the end of year -2—two years ahead of the bankruptcy filing for treated firms. Hourly wages are inflated to 2020 NOK. Numbers in brackets are standard deviations.

5 Effects of Firm Bankruptcies

In this section, we discuss the results from the estimated impact of bankruptcy filings on labor market and fertility outcomes. We also conduct a simple difference-in-difference estimation exercise to verify our findings, comparing average outcomes before and after the event, where we also discuss magnitudes of effects, and we conclude the section with a series of additional robustness checks on omitted trends, sample and bankruptcy event choices, and heterogeneous treatment effects.

5.1 Labor market outcomes

As an initial analysis into how labor market outcomes evolve before and after bankruptcy, Figure 6 compares the means over time for the men exposed to bankruptcies and their matched brothers. Note that these are sample means that do not account for any control variables. Even here, there is a clear divergence in outcomes after the bankruptcy event. Men experiencing a bankruptcy are substantially more likely to be unemployed, experience an earnings loss, have lower hourly pay and a dip in total working hours.

Next, Figure 7 displays the estimated coefficients from Equation (1) for each labor market outcome. Recall that this estimates the impact of the bankruptcy conditioning on a full set of individual and year * cohort fixed effects. Panel A depicts the impact on the individual having a valid record in the November file of the employer-employee register; this means having non-zero pay and non-zero contracted hours.[19] Bankruptcy is associated with a large decrease in employment, where individuals working in bankrupt firms are significantly more likely to be without a job compared to their siblings, and the effect is remarkably persistent. This finding is confirmed in Panel B, which shows a dramatic spike in unemployment probability during the year of bankruptcy. In addition to declining total earnings from work (Panel G), we also see declining after-tax income from all sources, including public transfers, in Panel H. This confirms that while the Norwegian social security system is generous, men facing a bankrupt employer do suffer long-term impacts on their after-transfer earnings in the region of a 8% decline.

An advantage of our setting is that we are able to analyze a long pre-treatment window. While the evolution of being registered for a November job is similar between treated and control brothers prior to time zero, we note that being registered as unemployed at any time in the year (Panel B), and earnings (Panels G and H) show differences between treated and control brothers already around three years before the bankruptcy. Brothers are on similar labor market paths and therefore comparable between four to seven years prior to the bankruptcy event (with the exception of log earnings which is also lower at year -4). This confirms the classic finding by Jacobson, LaLonde, and Sullivan (1993) and also the newer findings of Dustmann and Meghir (2005), who show that the impacts of a mass layoff or firm closure respectively are evident up to three years before the event. The key question is whether these patterns are indicative of worker heterogeneity in earnings growth that would have happened absent the bankruptcy event, or the impacts of a firm in distress leading up to a bankruptcy filing. If bankruptcy impacts are only evident at $t = 0$ and declining trends prior to this period are due to worker heterogeneity, then any difference in workers staying or leaving the firm should not be especially evident in the years leading up to bankruptcy.[20] A firm in distress would have slower wage growth, offer fewer paid hours, and lay off less productive workers

[19] We use November to avoid seasonal fluctuations in the summer months and around Christmas. This variable is also the basis for our sampling in the event study: the sample begins with the job held in November two years prior to the firm filing for bankruptcy.

[20] This is not to say that low productivity workers are not more likely to be laid off, but that this higher probability of lay off should not be concentrated in the years prior to bankruptcy if the bankruptcy event only has impacts at time zero.

in the years leading up to the bankruptcy filing. Similar to Jacobson, LaLonde, and Sullivan (1993), we argue that the latter interpretation is the appropriate one.

To show this, we investigate the evolution of hours worked and hourly pay in the lead up to bankruptcy filing for men exposed and unexposed to bankruptcy, comparing men still working at the firm or elsewhere (the "conditional sample"), and those who are no longer working at the firm and have found no replacement employment (adding these yields the "unconditional sample"). Using a conditional sample of men still in work, Panel C shows a substantial decline in the number of hours worked for treated men, relative to men still working at their control firm, as well as a decline in hourly pay (Panel E). These are important findings because the sample only consists of men actually working, whether at the bankruptcy firm or at a different firm. Even among this group of positively selected men, there is a declining trend in labor market outcomes for the treatment sample in the few years leading up to bankruptcy. This is consistent with a firm-specific trend. Next, we predict these outcomes for those men with missing data (the unemployed) - yielding an unconditional sample. Adding these men generates impacts that are *more negative* (Panels D and F), consistent with our argument that the sample in Panels C and E is one of positively selected men. See also the means in Figure 6. Thus, the to-be-bankrupt firm lays off less productive workers first and especially in the few years leading up to bankruptcy. This is consistent with the aggregate declining prior trend in working hours and hourly wage in Figure 7 explained by a firm-specific trend due to distress in the years prior to bankruptcy.

To follow the language of Jacobson, LaLonde, and Sullivan (1993), we find that bankruptcies in this context have effects over a number of years: they are associated with a "dip" in earnings from three years prior until year 0, an additional "drop" in the year of bankruptcy filing, and limited "recovery" in long run earnings. This pattern will have consequences for the timing of fertility and family outcomes, which we discuss below.

We also check the impact on an alternative outcome, registering for temporary and permanent disability, a preferred way of claiming benefits as a result of long-run unemployment. We see some increase in the average uptake of disability benefits, but this impact is not statistically significant at conventional levels (Panel I).

5.2 Fertility outcomes

In Section 3, we showed that men with a lower earnings rank are more likely to be childless and less likely to be partnered. However, these correlations could be driven by other factors for which we lack data. We use bankruptcies as a setting for an exogenous shock to earnings and other labor market outcomes. Having established these labor market impacts to be large and significant, we now analyze the impacts of a firm's bankruptcy filing on men's family outcomes.

In Figures 8 (means) and 9 (regression coefficients), we confirm the descriptive patterns in the bankruptcy setting. In particular, there is a divergence in average fertility and family outcomes after the bankruptcy event, which is especially evident for single status and total number of children. Turning to the estimated coefficients that account for individual and cohort*year fixed effects, we find that the probability of being unpartnered increases significantly following exposure to a firm bankruptcy, reaching a peak of 3.2 percentage points in year 4 (Panel A). The effect on single status somewhat lessens but does not dissipate over the 15 year period that we follow the men in our sample.

Turning to fertility, men exposed to the bankruptcy event are less likely to experience the birth of a child by 1.1-1.7 percentage points per year for at least six years following the event (Panel B; the impacts in later years are also significantly negative). The effect on experiencing a first birth - transitioning out of childlessness - makes up more than half of the effect (Panel C), consistent with the descriptive evidence from Section 3 that the extensive margin of fertility is the main driver of aggregate changes in fertility inequality. Thus, perhaps unsurprisingly, the probability of being childless increases significantly and remains positive for up to ten years after the bankruptcy event.

The impact on higher parity births is more muted and somewhat delayed compared to first births, being only statistically significant in years 5 and 6 (Panel D). Finally, the effect on total fertility is negative and grows over time (Panel F). The decline in total fertility is remarkably persistent and does not recover during our sample window of 15 years. These results resonate with the findings of Salvanes, Willage, and Willen (2021), who show that job loss impacts the fertility of men in the early career age group (in their case, 20-35) the most.

Considering the timing of impacts, the earliest effects are seen for flow variables: men experiencing a bankruptcy event are more likely to be single in a given year and less likely to experience the birth of a first child. These impacts are evident one year before the bankruptcy filing, thus two years after the impacts on annual unemployment first appear (see Figure 7). This makes intuitive sense as partner finding and fertility intentions take some time to materialise. Impacts on stock variables (childlessness status and total fertility) start to take shape at year 2. Taken together, we remark the following: the persistent negative effect on total fertility stems from "missed births" in the initial few years after job loss, and specifically first births, that are not compensated for in later life. The missed births themselves most likely stem from the reduced rate of partnering in the initial years around the time of the event.

Is partnering the key channel through which fertility outcomes materialize? To answer this question, we split the sample into those who are single at the time of sampling ($t-2$), and those that are in a couple at that time. We then estimate impacts separately for these two groups. Appendix Figure A.12 shows means and Figure A.13 the coefficients. Interestingly, we find that fertility declines for both by a similar amount: both (ex ante) partnered and unpartnered men are less likely to have a child after experiencing a bankruptcy filing. Although the pre-bankruptcy means of outcomes are naturally different between singles and couples, finding significant impacts for these groups shows that both channels are important: experiencing a bankruptcy reduces the ability to find a partner and have a child, but also reduces the likelihood of existing couples having a child together.

5.3 Difference-in-difference estimates

We verify that the patterns are robust to an alternative estimation strategy, namely a difference-in-difference regression that compares average outcomes before and after bankruptcy, collapsing the data to a pre- and post-period. This approach overcomes concerns about serial correlation in difference-in-difference approaches (Bertrand, Duflo, and Mullainathan 2004). It also allows a more intuitive measure of magnitudes. We discussed the fact that the impact of a bankruptcy on workers' labor market outcomes materialises over a number of years, beginning with a dip a few years before year 0 and then a decline in the year of the bankruptcy filing. Given this pattern, estimating a clean magnitude of the impact of a bankruptcy in this setting is not trivial. We opt to take a "donut" approach, removing years zero and -1 from the regression, thus comparing outcomes in years -7 to -2 (the "pre period") with outcomes in years 1 to 15 (the "post periods"; see column (1) in Table 3). In a second version (column (2) of the same table), we estimate an extended donut also omitting years -2 and -3 (thus fully removing the earnings "dip" years).

We find significantly lower income and higher unemployment probability in the post-bankruptcy period, as well as lower total fertility, a lower probability of having a child, and higher likelihood of still being childless. Using the baseline or extended donut does not make a substantive difference to the estimated impacts of bankruptcy in the difference-in-difference specification. These pre- and post-comparisons show clear and significant differences in family and labor market outcomes between treated and non-treated brothers.

Turning to magnitudes, we scale these effects by the post-period mean in the control group (which we interpret as a valid counterfactual for the treatment group had bankruptcy not occurred). We estimate an overall reduction in earnings of between 9-11.6% that is smaller than the comparable estimate from Jacobson, LaLonde, and Sullivan (1993) of a loss of 25%, though they follow individuals for a shorter period of five years after job loss. The probability of single status increases by 6.6% and childlessness by 4.4% of baseline, while the probability of fathering a child and total fertility decline by 10.1% and 2.4% respectively.[21,22] The comparability of the fertility magnitudes underlines the plausibility of our story, namely that following an inability to find a partner and "missing" the birth(s) of a child (children), permanent childlessness increases and total fertility never recovers.

Taken together, these results amplify the implications of our main findings. Men who face a negative labor market shock between the ages of 25 and 35 are less likely to have a child and to be partnered, and these effects remain 15 years after the shock, with very little recovery. Considering this along with our descriptive results on the cross-sectional inequality in family life across the earnings distribution, this uncovers an important connection between labor market prospects and

[21] These are compared to counterfactual means in the control group in the post-event period of 0.41, 0.07, 0.28 and 1.5, respectively.

[22] We discussed in Footnote 8 the two sources of data available for single status. If we restrict the bankruptcy sample to the 2004-2020 period in order to make use of the "rich" data on single status that would account for childless cohabitations, results are similar. Specifically, the coefficient of treat#post for single status is 0.021 (0.0075) and scaling this by the post period mean for the control group we compute an effect of 0.021/0.297*100=7.1%.

men's access to family life. In Section 6, we provide suggestive evidence that the fertility "penalty" to job loss has increased over time, and so put forward a potential explanation for why inequality in fertility has grown over the last three decades.

Figure 6: Sibling mean comparisons before and after firm bankruptcies, labor market outcomes.)

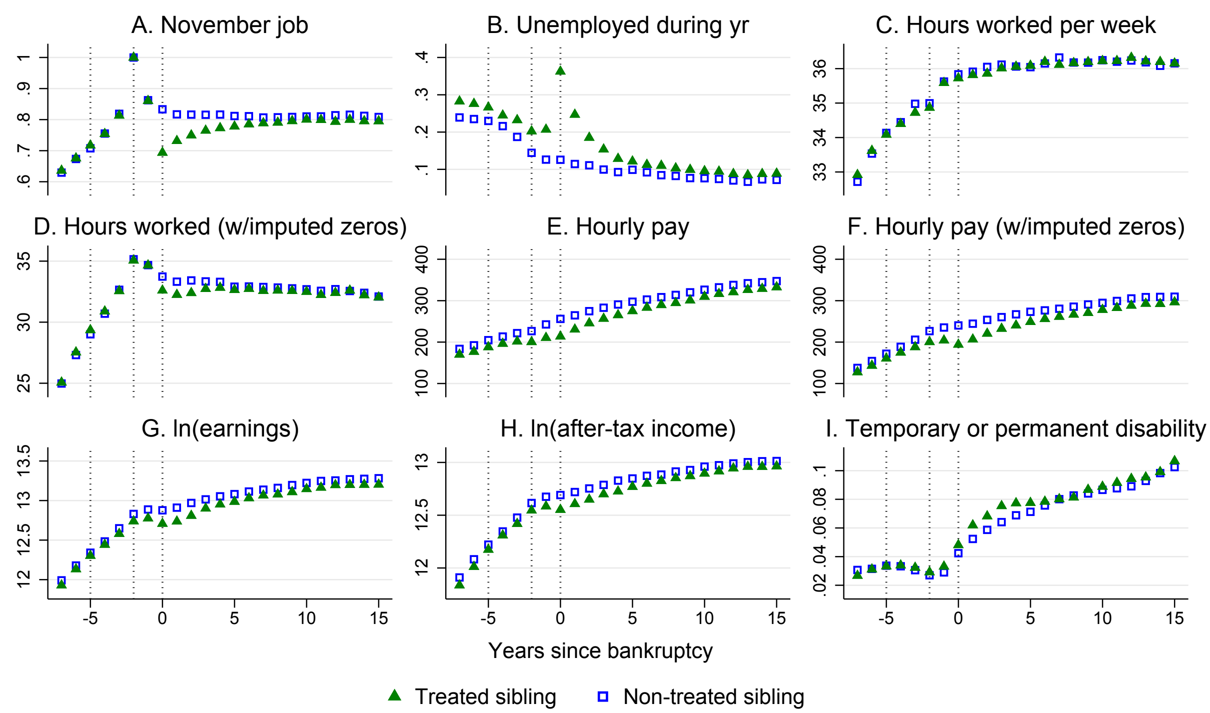

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Sample of treated siblings consists of Norwegian-born men who in year -2 were employed in a firm filing for bankruptcy two years later and age 25-35 in the year of the event, while non-treated siblings in year -2 held a job with an employer that did not file for bankruptcy during the observation period. Samples are restricted to families with both treated and non-treated siblings. Samples in Panels C and E are further restricted to those with a job in November, while Panels D and F impute hours and pay for those without a job. Full sample observation counts are 267 292 in the treatment group and 317 579 in the control group, but may be lower for some outcomes with missing data in certain years. Wages and earnings are inflated to 2020 NOK.

Figure 7: Effects of firm bankruptcies on labor market outcomes.

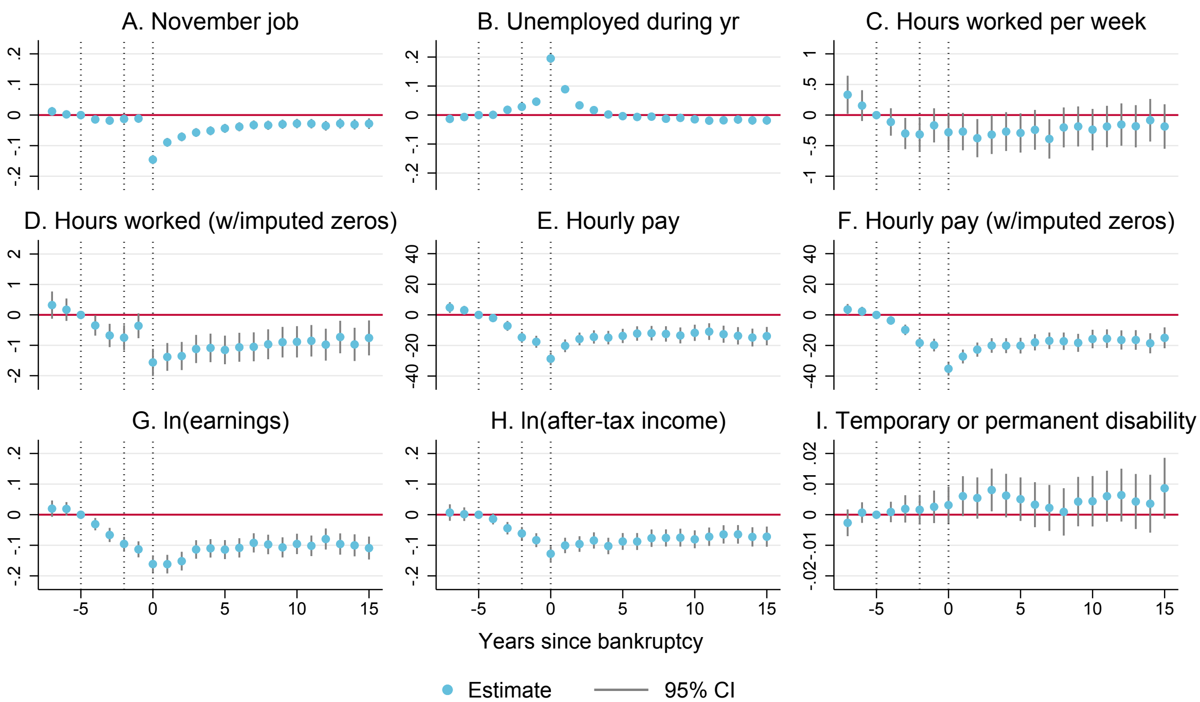

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Sample of treated siblings consists of Norwegian-born men who in year -2 were employed in a firm filing for bankruptcy two years later and age 25-35 in the year of the event, while non-treated siblings in year -2 held a job with an employer that did not file for bankruptcy during the observation period. Samples are restricted to families with both treated and non-treated siblings. Samples in Panels C and E are further restricted to those with a job in November, while Panels D and F impute hours and pay for those without a job. Observation counts are 267 292 in the treatment group and 317 579 in the control group.

Figure 8: Sibling mean comparisons before and after firm bankruptcies, fertility outcomes.

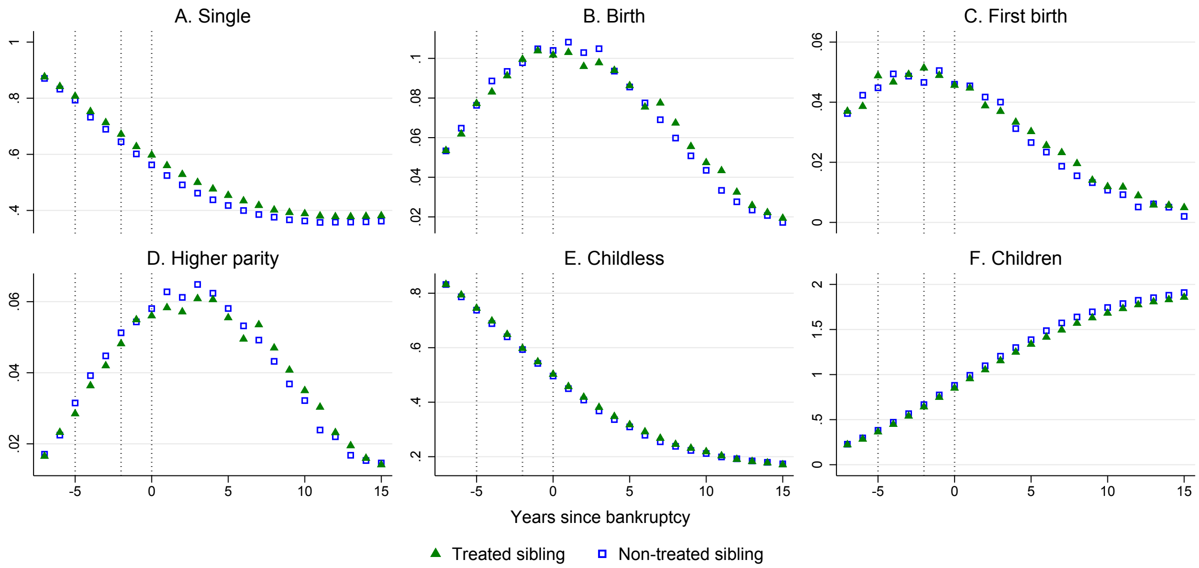

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Sample of treated siblings consists of Norwegian-born men who in year -2 were employed in a firm filing for bankruptcy two years later and age 25-35 in the year of the event, while non-treated siblings in year -2 held a job with an employer that did not file for bankruptcy during the observation period. Samples are restricted to families with both treated and non-treated siblings.

Figure 9: Effects of firm bankruptcies on fertility outcomes.

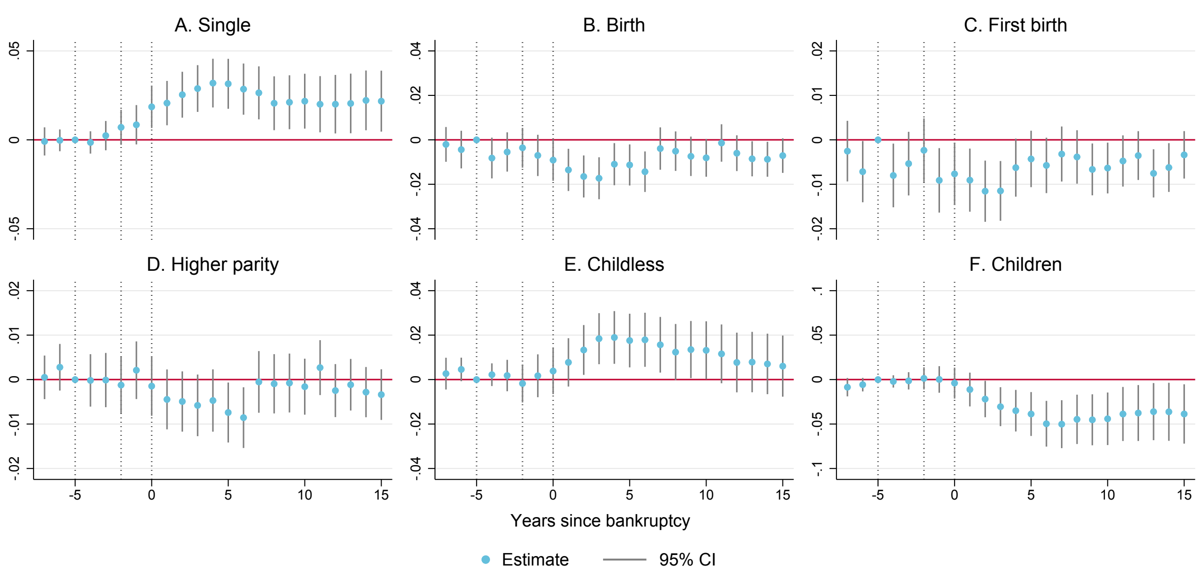

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Scatter points show the estimates of β_t from the estimating equation. See text and notes to 8 for a description of the samples.

Table 3: Difference in difference regression estimates

Outcome	(1) Coefficient (omit yrs -1 and 0)	(2) Coefficient (omit yrs -3 to 0)
Labor market outcomes		
November job	-0.044***	-0.046***
	(0.005)	(0.005)
Unemployed during yr	-0.001	0.009**
	(0.004)	(0.004)
Hours per week	-0.897***	-1.114***
	(0.167)	(0.200)
Hourly pay	-14.079***	-19.378***
	(2.072)	(2.441)
ln(earnings)	-0.090***	-0.116***
	(0.011)	(0.013)
ln(after-tax income)	-0.066***	-0.084***
	(0.009)	(0.011)
Fertility outcomes		
Single	0.025***	0.027***
	(0.006)	(0.006)
Birth	-0.006***	-0.007***
	(0.002)	(0.002)
First birth	-0.002**	-0.002
	(0.001)	(0.001)
Higher parity birth	-0.004***	-0.005***
	(0.001)	(0.001)
Childless	0.013**	0.012***
	(0.005)	(0.005)
Children	-0.035***	-0.035***
	(0.011)	(0.012)
Observations[1]	526 798	468 547
Individuals	29 204	29 201

Notes: This table displays the regression coefficients from a series of regressions comparing outcomes pre and post bankruptcy, between treated and control brothers. Column 1 compares outcomes in years -7 until -2 with outcomes in years 1 to 15, while Column 2 omits years -2 and -3, comparing outcomes in years -7 to -4 with years 1 to 15.
* denotes p-value<0.1, ** denotes p-value<0.05 and *** denotes p-value<0.01.
[1]Observation count is for fertility outcomes, count may be smaller for other outcomes because of missing data (yrs w/o data) and/or log zero problem. For comparison, count for ln(earnings) in column 1 is 491 620.

5.4 Linking the descriptive and causal evidence

The striking facts on male fertility discussed in Section 3 motivated an analysis that attempts to causally estimate the response of men's family outcomes to exogenous variation in income, and so investigate whether the descriptive relationships reflect a causal relationship or the impact of other correlates. Our results show that, in our micro setting of prime-age men facing bankruptcies, there is a strong causal link between men's labor market performance and their partnering and fertility outcomes. Next, we conduct a back-of-the-envelope exercise to estimate the share of the broad descriptive relationship evidenced in Section 3 that could be explained by a causal earnings-fertility relationship along the lines estimated in this Section.

To do this, we first make the samples comparable across the two exercises. In particular, as the event study estimates focus on men who experience a bankruptcy between ages 25 and 35 and are therefore aged 18-50 in the sample with a mean age of 33, we show the descriptive patterns of fertility as a function of real earnings between ages 21 to 40 (rather than the rank of lifetime earnings). To better match birth cohorts of the event study sample, we also restrict the descriptive patterns to those born between 1971 and 1980. Figure 10, Panels A and D, display binned scatter plots of total fertility and childlessness against real earnings, showing similar non-linear patterns as those for the most recent cohorts in Figure 1. These patterns remain highly non-linear when plotted against log earnings (Panels B and E), but when we trim the data for the bottom and top 5 percentiles of log earnings, Panels C and F show that the relationships between childlessness and children by age 40 and log earnings are well approximated by linear regressions. Estimating these regressions for these birth cohorts, we find that a one log point increase in earnings is correlated with a reduction in the probability of childlessness of 23.2 percentage points, and with having 0.62 more children at age 40.

We next turn to the estimated effect of a bankruptcy on labor market and fertility outcomes in the difference-in-difference specification. Referring to column (2) in Table 3 we note that a bankruptcy reduces earnings by 0.116 log points and the number of children by 0.035, while raising the likelihood of childlessness by 1.2 percentage points. Scaling this up to 1 unit of log earnings yields magnitudes of 0.3 children and 10.0 percentage points of childlessness.

While we do not put forth that a formal analysis using bankruptcies as an instrumental variable for earnings would satisfy the exclusion restriction, as bankruptcies are likely to affect multiple outcomes including time use and self-esteem, we think it is nevertheless useful to bring together these two sets of estimates for an informal calculation of the share of the descriptive evidence that is likely to be causal. Scaling these two sets of effects indicates that around 48% (0.3 / 0.62 * 100) of the descriptive relationship between earnings and total fertility and 43% (10.0 / 23.2 * 100) of the comparable relationship between earnings and childlessness can be explained by a causal relation. This emphasises the plausibility of our estimates.

Figure 10: Binned scatter plots, fertility outcomes age 40 and real earnings at ages 21-40.

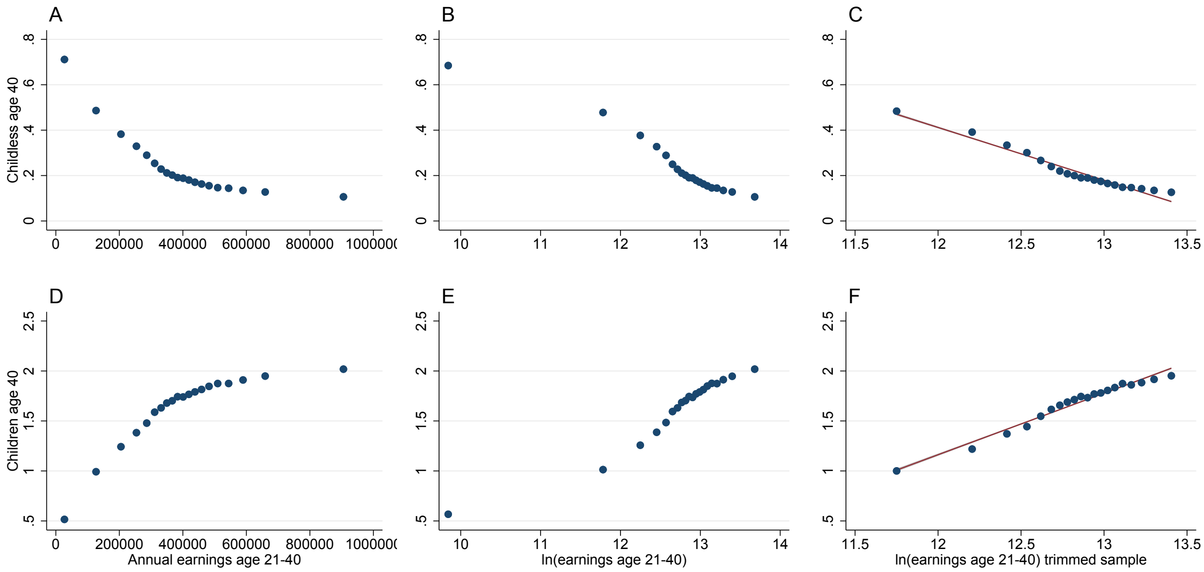

Notes: Sample consists of men born between 1971 and 1980. Earnings are inflated to 2020 NOK. Sample in Panels C and F is restricted to 5th to 95th percentile range of ln(earnings). Slopes (se) of regression line are -0.232 (0.002) in Panel C and 0.615 (0.006) in Panel F. Shaded area around regression line depicts 95 percent CI of prediction. Observation counts are 257 526 in Panels A and D, 256 496 in Panels B and E, and 230 848 in Panels C and F.

5.5 Robustness Checks

We conduct important checks on omitted trends, sample and bankruptcy event choices, and heterogeneous treatment effects in this Section.

Omitted trends The findings on the impact of bankruptcies on labor and fertility outcomes are identified on the assumption that there are no omitted trends that drive the path of labor and fertility outcomes for treated men. We explore several checks on this assumption. In our main specification, we sample individuals employed at a to-be-bankrupt firm two years prior to bankruptcy, similar to one of the specifications in Dustmann and Meghir (2005). This may lead to some selection on outcome variables, and selecting the sample in an earlier year will improve the exogeneity of this sample (with respect to bankruptcy) but also increase measurement error with less precise treatment. We explore how these samples differ by conducting an event study-type analysis of outcomes over time for our main treatment sample, our control sample, as well as a few alternative samples: individuals employed at the to-be-bankrupt firm five years prior to bankruptcy, four years prior, and three years prior, as well as individuals employed two years prior but not satisfying the additional condition of having a same-sex sibling in the control sample. The time paths of our main outcomes for these different samples are shown in Figure 11.

The time paths are surprisingly similar across all samples. There are notable deviations from trend for unemployment in the year following when we restrict individuals to be employed: for example, there is a spike in unemployment at t-4 in the sample whose last year of employment at the firm is t-5. This is a direct result of this definition and to be expected. More remarkably, the time paths of family outcomes - partnering status, births and total children - are surprisingly similar across all groups. This indicates that our choices of treatment and control samples do not induce a large amount of selection on trends in outcomes.

Figure 11: Evolution of outcomes over time for different samples.

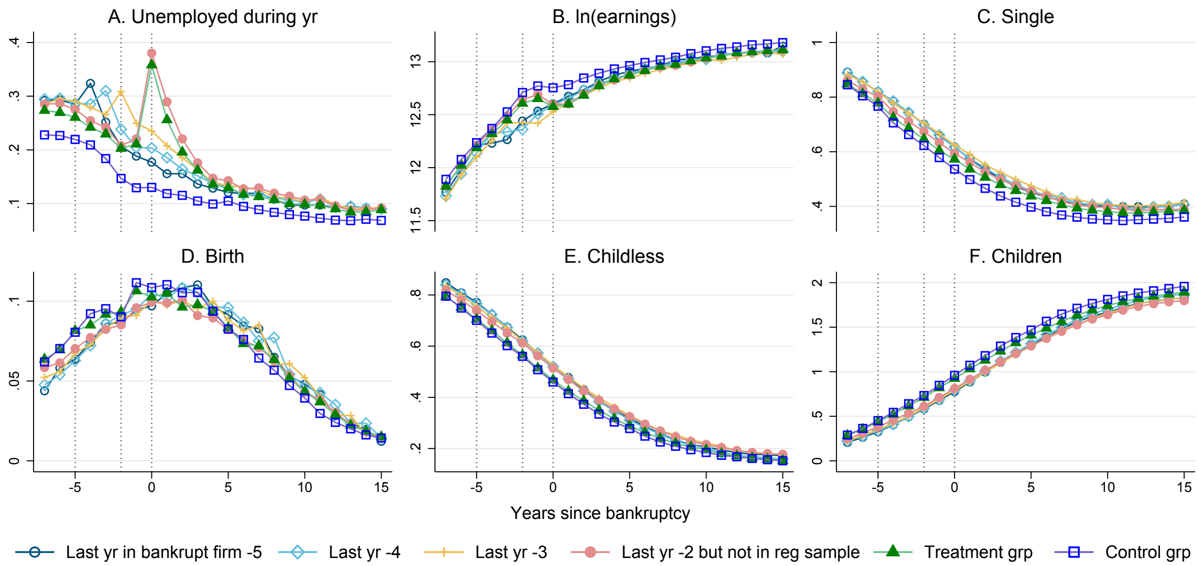

Notes: Samples consist of men age 25-35 at time 0, separated by the last year of employment at the firm filing for bankruptcy (at time 0). For completeness, the figure adds the time paths for the relevant outcomes for treatment and control groups depicted in Figures 6 and 8.

Although the time paths are reassuringly similar, we conduct an additional robustness check where we estimate our main specification but sample individuals a year earlier, at t-3. This is expected to change the sample composition: while the sample may be more exogenous in the sense that there is less selection into (or out of) a firm that will eventually be bankrupt, there will also be more measurement error in treatment because fewer of these individuals will actually experience the bankruptcy event that arises in three years' time.

Figure A.15 show the results (Figure A.14 in the Appendix shows the evolution of means between the two samples). Our main findings on labor market outcomes, marital status and total fertility are robust to this alternative sample definition, though smaller in magnitude. The impact on births is less marked here, with coefficient estimates that are not statistically significant. We interpret this as showing us the impact of more measurement error when choosing a less precise treatment sample.

An alternative sampling strategy for the control group follows Fadlon and Nielsen (2019) and Salvanes, Willage, and Willen (2021). Specifically, we define the control group to consist of men who *do* experience a firm bankruptcy but in later years (2006-2015). We include these men as a control group when they are ages 25-35 in the years between 1995-2005. (Note that this means that the control group are much older when they experience their bankruptcy - the median is 42 in our data). The treatment sample is defined on bankruptcies between 1995-2005, but neither control nor treatment sample are now restricted to have brothers. Appendix Figure A.16 shows the estimates, which look reasonably similar to our baseline estimates in the early period before the control sample eventually experiences their bankruptcy event - particularly on stock fertility outcomes.

Another way of accounting for possible omitted trends is to include time-varying fixed effects. Given our treatment is at the individual-year level, time-varying fixed effects to account for individual heterogeneity need to be at a level above the individual. The lowest possible level to consider is siblings (pairs of brothers). We re-estimate our model adding family * year fixed effects. These account for any family-specific characteristics that may vary over time, such as common trends in family outcomes specific to siblings. One example is that brothers from a large family may have a steeper positive trend in total fertility than brothers from small families. Figure A.17 shows that the estimates are essentially unchanged.

Alternative sampling and event choices Next, we show that our results are robust to alternative definitions of bankruptcies, and altering the bankruptcy sample in various ways. First, we examine whether our results are robust to an alternative definition of workplace closure, turning to establishments and using any event where the number of employees at the establishment drops to zero and does not recover. To minimise false shutdowns due to mergers or acquisitions, we override the shutdown event if two thirds or more of last year's workforce work at the same establishment at the end of the shutdown year. The approach is in line with that used in prior studies, such as Rege, Telle, and Votruba (2007), Huttunen, Møen, and Salvanes (2011), and Salvanes, Willage, and Willen (2021) and yields a more broad definition of workplace closure.

Figure A.18 in the Appendix shows mean outcomes over time for the two comparison groups. Figure A.19 shows that, although our main estimated effects on unemployment and earnings persist here, they are smaller in magnitude than those in Figure 7. Consistent with the smaller effects on economic outcomes, the estimated effects on single status and fertility are also attenuated when compared to those from bankruptcies (though are statistically significant in the period shortly after establishment shutdown). However, there are two reasons why we prefer bankruptcies over this broader measure of firm shutdowns. First, although we minimise false shutdowns, we may not be able to rule them out entirely, which can introduce measurement error. Second, the closure of an establishment likely represents a less abrupt change compared to a firm bankruptcy, and hence may be more easily anticipated. In this sense, we can expect some attenuation of estimated coefficients.

Our estimation sample relies on selecting individuals working at the treated firm two years prior to its bankruptcy. This is matched by a sibling sample working in a stable firm. However, this does not preclude that a bankruptcy was experienced by the treated sample in any year before or after -2 (a separate bankruptcy at another firm), or that the sibling experienced a bankruptcy in another year. As a robustness check we apply a more stringent criterion to our sample by restricting our treated sample to individuals who only experienced the bankruptcy of interest, and siblings who never experienced a bankruptcy. Figure A.20 in Appendix A.2 shows the estimates, which are not sensitive to this stricter sample restriction.

A second source of contamination arises when bankruptcies are sufficiently large to affect a significant share of the workforce in the local labor market, also influencing the labor market opportunities and the economic status of the control group not directly involved in a bankruptcy. To address this concern, we identified years where at least 1.5 percent of the municipality workforce was subject to a bankruptcy two years later and dropped any such event sequences from our estimation sample, reducing the sample by 11 percent. Figure A.21 in Appendix A.2 shows the estimates from this exercise, which again are not sensitive to the stricter sample restrictions.

A converse concern relates to small bankruptcies, where adverse life experiences of workers in small firms may be the direct cause of the negative economic performance of the soon-to-go-bankrupt firm, leading to reverse causality in the treatment group. A plausible check is to reestimate the model, dropping bankruptcies in small firms from the sample. Figure A.22 in the same Appendix show the estimates after we exclude bankruptcies in firms below the 10th percentile of the firm size distribution in the treatment group (i.e., firms with less than six employees in the base year). As the figure shows, results are not sensitive to this concern.

We also verify that our results are not sensitive to whether our sample is balanced or not. In our main estimation sample, we do not make the restriction that all included men are observed in all years. Here, we restrict the sample to those with bankruptcy years 1997-2005 (compared to the baseline 1995-2015) whose unemployment and fertility outcomes can be tracked for the full 23 years.[23] We find that in this restricted, smaller sample, coefficient estimates and patterns are very similar to those in the main estimates, but with expectedly wider confidence intervals (Figure A.24, with means reported in Figure A.23 in the Appendix).

Heterogeneous treatment effects An issue in staggered regression designs with two-way fixed effects is that estimates draw on already treated units as controls for units that are treated late in the sample period, rendering bias in estimates of counterfactual outcomes when there are heterogeneous treatment effects (see, e.g., Goodman-Bacon 2021, Callaway and SantAnna 2021, and Sun and Abraham 2021). In our setting we have individuals that are never treated, i.e., the siblings, and the mean comparisons of trajectories of treated and non-treated siblings (as in, e.g., Figure 6) do not suffer from this problem. Our estimates may nonetheless be subject to this type of bias if sample inclusion of already treated individuals influence estimation of calendar year effects, which we condition on when estimating counterfactual trajectories.

To address this concern, we follow Cengiz, Dube, Lindner, and Zipperer (2019) and conduct a stacked event-by-event analysis. In this analysis we take each of the 21 bankruptcy years in our data and generate "clean" samples, i.e., excluding any other observations that have already

[23]Civil status is first available in 1991, however, while data on earnings end in 2019.

been treated, for each of the post-event trajectory years. We then run separate regressions for each combination of bankruptcy and trajectory year and aggregate the estimates. We present the results in Figure A.25, where we see that the point estimates are similar to those from the baseline approach but that we lose precision in using the smaller stacked samples (where underlying point estimates on average draw on only 1/21 of available treatment observations). Although there are some detectable differences in estimates of effects on log earnings, the important take-away from this exercise is that there is no indication that sample inclusion of already treated observations renders bias in estimates of effects of bankruptcy on family outcomes.

6 The Changing Relationship between Labor Market Outcomes and Fertility

The correlations described in Section 3 show important inequalities in male fertility that have increased over time. Examining the impact of bankruptcies on male fertility in Section 5, we document similar patterns in a more causal way. However, this analysis does not speak to the change in this relationship over time. In this section, we bring the two exercises together. We investigate whether the relationship between negative labor outcomes and fertility has changed over time, and whether this can plausibly explain the important facts on growing inequality in family outcomes.

There are three ways in which growing inequality in fertility could be explained by changes in the labor market, which we together refer to as a "fertility penalty". The first evolves as follows: with the increasing economic independence of women, and rising living (and child rearing costs), men's earnings may have become a more important determinant of partner finding. In this sense, experiencing a negative event that leads to earnings loss may have a larger impact on fertility now compared to earlier, even if the earnings loss is the same.

A second mechanism can occur via changes in the labor market itself. In particular, the same event (bankruptcy, or job loss) may carry larger earnings penalties now, compared to in earlier years, for example because of stagnating job skills or education that would otherwise improve the resilience to such shocks.[24] A third mechanism relates more broadly to the other impacts of job loss that are not purely economic: changes in identity and self-esteem. These factors will also affect men's ability to find a partner and form a family. To the extent that these impacts have changed or worsened over time, this can also explain a growing fertility penalty over time.

These three channels could drive the growing inequality in men's fertility outcomes over time. In the first mechanism, earnings losses are likely to be constant in response to the same event over time. Of course, growing earnings impacts do not preclude a higher earnings threshold for finding a partner, but we can exclude the second mechanism if we find no evidence of changing impacts of shocks on earnings. The third mechanism is more difficult to find evidence for, but we suggestively

[24] Alternatively, there could be a larger number of negative events in recent years, but our dataset shows no growth in the number of bankruptcies over time.

analyze changes in the rate of re-employment as a proxy for impacts on identity and self-esteem.

To shed light on these important questions, we take two approaches. In the first approach, we use the bankruptcy setting and compare impacts between an early and late period. In the second approach, we explore correlations between previous unemployment status and current fertility, earnings and re-employment across calendar years.

6.1 The Changing Impact of Bankruptcies over Time

In this first analysis we use the same bankruptcy setup as before and split the data by the event year into an early period (1995-2005) and a late period (2006-2015). If the penalty to job loss has grown over time, we would expect larger coefficient magnitudes for the late period, compared to the early period. The analysis will inevitably lack statistical power: with full individual, age and cohort fixed effects, adding interaction terms for each time period will challenge the data (and especially for the last time periods of the late period, where there are the fewest observations). Nevertheless, the patterns are striking. We see in Figure 12 that most coefficients are larger in magnitude for the late period, after bankruptcy filings.[25] Men experiencing bankruptcies in the latter half of our study period experience higher rates of childlessness and lower fertility compared to those in the first half of the study period. Strikingly, the earnings impacts also appear to be larger in the late period, and re-employment prospects seem worse (Panels A and B). This suggests that experiencing a negative labor market event carries a greater earnings penalty in recent years, compared to earlier. In this sense, at least part of the growing inequality in fertility could be explained by a more challenging labor market that fails to shield vulnerable workers.

6.2 The Changing Fertility Penalty to Unemployment

In order to obtain more statistical power, our second analysis takes a slightly different approach. Using cross-sectional population data for the period 1990-2020 for individuals aged 25-35 to match the event study sample, we regress several outcomes on individual, lagged unemployment status while controlling for years of education, potential labor market experience and its squared term, and municipality fixed effects, akin to a Mincer regression. We estimate this regression with flexible interactions to allow for the coefficient of lagged unemployment status to vary with the year of observation. In particular, we estimate:

$$\begin{aligned} Y_{i,k,t} = \alpha + \sum_{t=1990}^{t=2020} \beta_t Year_{i,t} * Unemp_{i,k,t-1} + \gamma Exper_i + \tau Exper_i^2 \\ + \lambda Educ_i + \kappa_{i,k} + \theta_t Year_{i,t} + \eta_{i,k,t}, \end{aligned} \tag{2}$$

where we consider three different outcomes $Y_{i,k,t}$: $Birth_{i,k,t}$ indicates individual i having a child in year t living in municipality k, $Find_{i,k,t}$ measures the rate of job finding in the year

[25]Note that we do not yet observe in the data the 14 and 15 year impact of bankruptcies that occurred in 2006.

Figure 12: Effects of firm bankruptcies on labor market outcomes in different periods.

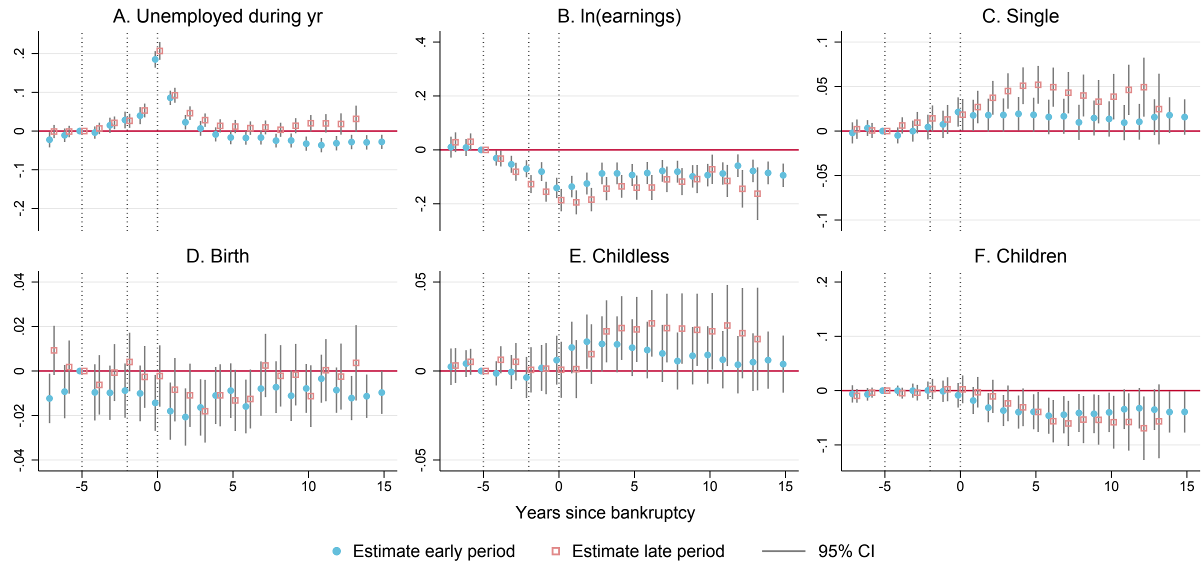

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). The sample is the same as in Figure 9 but we split the sample into an early period (those with event years 1995-2005) and a late period (event years 2006-2015).

following job loss and $Earn_{i,k,t}$ are annual earnings. Our right-hand side variables are: $Unemp_{i,k,t-1}$ indicates individual i's unemployment status in the previous year, $Exper$, $Exper^2$ and $Educ$ are the individual's working experience, its quadratic, and their years of education, $\kappa_{i,k}$ are fixed effects for municipality of residence and $Year_{i,t}$ is calendar year. We focus on the set of coefficients β_t, which capture the relationship between outcomes this year and last year's unemployment status by calendar year.

In Figure 13, the top panels A-C show mean birth rates and the bottom panels depict the coefficients on lagged unemployment from this regression for any birth, along with similar estimates from regressions restricting the sample to first births and higher parity births.

The top panels show that fertility has been declining over time, with birth rates falling over the sample period. They also show a widening gap over time between those unemployed and those not. Focusing on the regression coefficients, Panel A shows that the relationship between unemployment and birth has become more negative over time: while being unemployed is associated with a higher probability of not fathering a child, this probability is larger in more recent years. Panel B shows that this effect is mostly driven by first births: unemployed men are less likely to transition out of childlessness the following year, and this probability has increased over time. Again, we confirm the importance of the extensive margin of family formation in explaining inequality in fertility: men suffering negative labor market events are less likely to form their first family. The trend line is statistically significant and confirms that the coefficients are declining over time. Panel C shows that the relationship for higher parity births is also negative, but with a less clear downward trend over time.

To shed light on mechanisms, the same figure shows estimates for re-employment (Panel D) and annual earnings (Panel E) as outcomes. For both panels, there appears to be a downward trend in coefficients across calendar years. Although the pattern is modest, simple regressions indicate a statistically significant, negative trend in coefficients for both outcomes. Thus, there is evidence to suggest that job loss has carried a higher earnings penalty in the latter half of this period, but also that re-employment prospects have deteriorated. In terms of partner finding, both are likely to be crucial: while lower earnings will have made men less attractive all else equal, increasing struggles to find a new job will affect their self-esteem, thus further bolstering the negative impacts of job loss on partner finding and family formation.

An important alternative explanation for the patterns in this figure is that there have been changes to the composition of the unemployed over time. For example, with declining unemployment over time, there may be an increasingly select sample of lower quality men who also have worse fertility outcomes, ex ante. To check this hypothesis we conduct two exercises: first, we track unemployment rates and the average education level of those unemployed, over time. Second, we show how the coefficient estimates in Figure 13 vary with the addition of controls related to this type of selection, such as education, BMI and IQ. Figure A.26 in the Appendix does not point to any strong evidence for changes in sample selection, as the difference in average education between the unemployed and the employed in our estimation sample has remained relatively stable over time. There is a slight increase in the average education difference after 2007 but this is not reflected in a change in the birth-unemployment coefficients around this year, which are mostly stable after this point.

Next, Figure A.27, also in the Appendix, illustrates that the largest change in estimated coefficients occurs with the addition of education controls, which account for around half of the estimated impact with minimal controls; however, the decrease in the magnitude of the coefficients over time is unaffected by the inclusion of the education controls, they only have a level effect. This is important, as it suggests that while education may mediate the size of the impact of unemployment on fertility, it does not explain the change in this impact over time. Interestingly, the addition of IQ and BMI controls does not make a meaningful change to estimated coefficients, showing that the education controls are well able to capture any individual differences that play a role in labor market outcomes. To sum up, while education is important, we do not find strong evidence that changes in selection over time are a primary explanation for the patterns seen in Figure 13.

These striking findings show that job loss carries a higher "fertility penalty" in recent years, consistent with the population patterns depicted in Section 3. Men experiencing poor labor market outcomes in recent years are more likely to be left behind in terms of family outcomes, and specifically having children. We find suggestive evidence that both earnings and re-employment have suffered more over time for the unemployed. Over and above any growing importance of earnings in partner finding, this shows that men are facing a tougher labor market with less recovery after bad events.

Figure 13: Unemployment, births, job finding and earnings over time

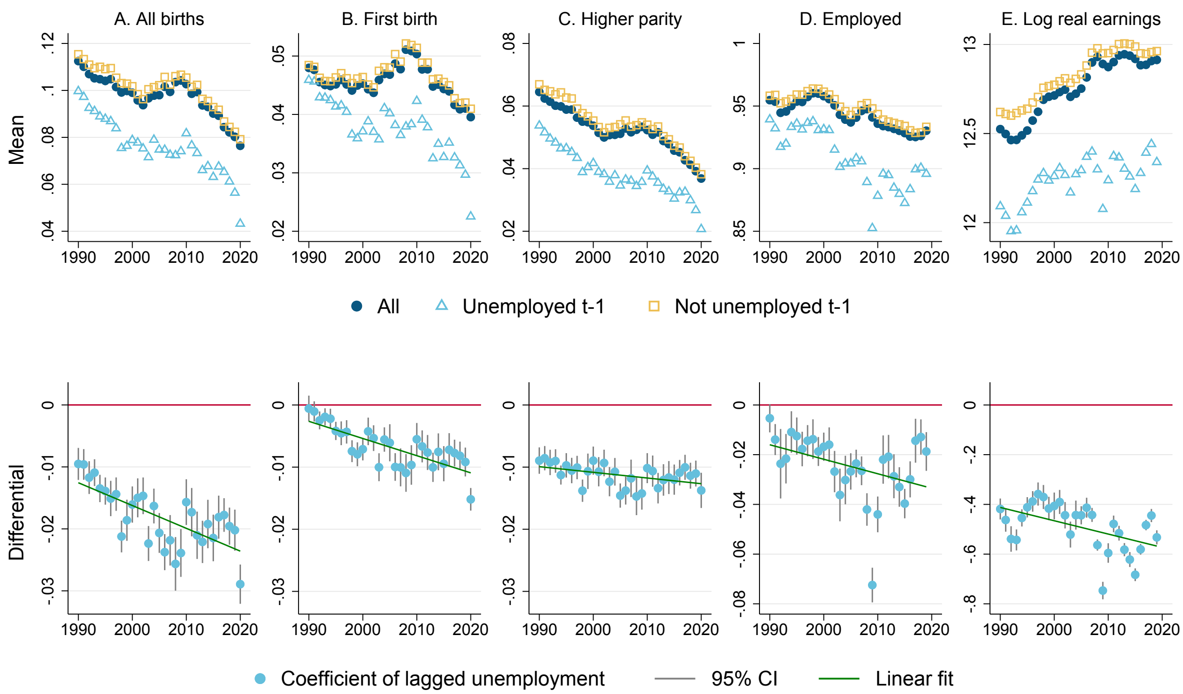

Notes: Scatter points in the top panels show fertility rates, current calendar year employment rates and log annual earnings of men by unemployment status of the prior calendar year, while the bottom panels show the estimated coefficient of individual unemployment status from a regression of these outcomes on registered unemployment the prior year. Regressions control for educational attainment, experience and its square, year of observation, and municipality fixed effects, and allows for the coefficient of lagged unemployment to vary by observation year. Standard errors are clustered by municipality. Sample consists of men age 25-35, sample period is 1990-2020. Observation count is 9 126 238 (fertility), 9 024 172 (employment) and 8 746 188 (log earnings). Mean birth rate is 0.098 and mean registered unemployment is 0.141. Slopes (standard errors) of regression lines are -0.00037 (0.00007) in Panel A, -0.00028 (0.00004) in Panel B, -0.00009 (0.00003) in Panel C, -0.00058 (0.00026) in Panel D and -0.0054 (0.0017) in Panel E.

Taken together with our findings from the bankruptcy analysis, a clear picture emerges that men's family outcomes are shaped by their labor market prospects. Job loss and its associated negative labor market outcomes lead to lower fertility, higher childlessness, and less partnering, with a penalty that has been growing over the last three decades, particularly for the extensive margin of family formation.

7 Conclusion

Using detailed administrative data from Norway, we document a remarkable increase in the inequality of male childlessness across the income distribution, with low-income men facing a "retreat from fertility". We further show that the poorest men are more likely to be single and that the income gradient in partnership formation has become steeper. These inequalities are primarily driven by the extensive margins of fertility and partnering. To investigate whether the labor market may causally explain these descriptive facts, we use bankruptcies to identify the effect of job loss on fertility. We note significant negative impacts of bankruptcies on employment, earnings, births, total fertility and partnering rates that do not recover for up to 15 years following the event. A simple calculation indicates that between 43%-48% of the descriptive earnings-fertility gradient may be driven by a causal relationship. We further show that the effect of bankruptcies on outcomes is stronger in more recent years, and that the relationship between unemployment and fertility has also become more negative over time, particularly for transitions out of childlessness. Worsening earnings and re-employment prospects seem to be mediating factors. A stronger "fertility penalty" to job loss has emerged in recent years, driven by a changing labor market that fails to shield low income workers with job insecurity.

Our work is distinct from the literature in two ways: by overcoming data limitations, we are able to make statements about the distribution of male fertility, and by including all men, rather than only men with partners, we can capture the total ramifications of job losses on male fertility, which often work through changes in partnering.

More broadly, the fertility changes we document show evidence for an emerging pattern in the marriage market: a type of "serial monogamy" (de la Croix and Mariani 2016), akin to a modern form of polygyny, where some men have children with multiple partners, while a significant portion of low income men do not participate in family formation. In this sense, low income men face wider consequences of stagnating earnings that reach beyond their labor market prospects. We note that this inequality may increase even further with the presence of a marriage premium for men, and possibly also a father premium, whereby earnings increase as a result of partnering (see Juhn and McCue 2017 for an overview and Kunze 2020 for recent evidence from Norway). Our results have wider societal ramifications that are not captured by focusing on earnings changes only: we document a dramatic shift in the distribution of new births in the population, which is likely to be accompanied by changes in child investments and quality of the next generation (Fagereng, Mogstad, and Rønning 2021).

References

Adda, J., C. Dustmann, and K. Stevens (2017). The career costs of children. *Journal of Political Economy 125*(2), 293–337.

Almås, I., A. Kotsadam, E. R. Moen, and K. Røed (2020). The economics of hypergamy. *Journal of Human Resources*.

Andrew, A. and A. Adams-Prassl (2022). Revealed beliefs and the marriage market return to education. *Mimeo*.

Anelli, M., O. Giuntella, and L. Stella (2019). Robots, labor markets, and family behavior. *IZA Discussion Paper Series* (12820).

Autor, D., D. Dorn, and G. Hanson (2019). When work disappears: Manufacturing decline and the falling marriage market value of young men. *American Economic Review: Insights 1*(2), 161–178.

Autor, D. and M. Wasserman (2013). Wayward sons: The emerging gender gap in labor markets and education. *Third Way Working Paper*.

Baudin, T., D. de la Croix, and P. E. Gobbi (2015). Fertility and childlessness in the United States. *The American Economic Review 105*(6), 1852–1882.

Bergsvik, J., A. Fauske, and R. K. Hart (2020). Effects of policy on fertility: A systematic review of (quasi) experiments. *Statistics Norway Discussion Paper No 922*.

Bertrand, M., E. Duflo, and S. Mullainathan (2004, 02). How Much Should We Trust Differences-In-Differences Estimates? *The Quarterly Journal of Economics 119*(1), 249–275.

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

Bhalotra, S., A. Venkataramani, and S. Walther (2022). Fertility and labor market responses to reductions in mortality. *Mimeo*.

Black, D. A., N. Kolesnikova, S. G. Sanders, and L. J. Taylor (2013). Are children "normal"? *Review of Economics and Statistics 95*(1), 21–33.

Bratsberg, B., O. Raaum, and K. Røed (2018). Job loss and immigrant labour market performance. *Economica 85*(337), 124–151.

Callaway, B. and P. H. SantAnna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Case, A. and A. Deaton (2020). *Deaths of Despair and the Future of Capitalism.* Princeton University Press.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics 134*(3), 1405–1454.

Chiappori, P.-A., S. Oreffice, and C. Quintana-Domeque (2012). Fatter attraction: Anthropometric and socioeconomic matching on the marriage market. *Journal of Political Economy 120*(4), 659–695.

Comolli, C. L., G. Neyer, G. Andersson, L. Dommermuth, P. Fallesen, M. Jalovaara, A. K. Jónsson, M. Kolk, and T. Lappegård (2020). Beyond the economic gaze: Childbearing during and after recessions in the Nordic countries. *European Journal of Population*, 1–48.

Craig, J. M., B. Diamond, and A. R. Piquero (2013). Marriage as an intervention in the lives of criminal offenders. In *Effective Interventions in the Lives of Criminal Offenders*, pp. 19 – 37. New York, NY: Springer New York.

de la Croix, D. and F. Mariani (2016). From polygyny to serial monogamy: A unified theory of marriage institutions. *Review of Economic Studies 82*, 565–607.

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

Doepke, M., A. Hannusch, F. Kindermann, and M. Tertilt (2022). The economics of fertility: A new era. *NBER Working Paper* (2).

Dustmann, C. and C. Meghir (2005). Wages, experience and seniority. *Review of Economic Studies 72*, 77–108.

Ellingsæter, A. L. (2006). The Norwegian childcare regime and its paradoxes. *Politicising parenthood in Scandinavia. Gender relations in welfare states*, 121–144.

Fadlon, I. and T. H. Nielsen (2019). Family health behaviors. *American Economic Review 109*(9), 3162–91.

Fagereng, A., M. Mogstad, and M. Rønning (2021). Why do wealthy parents have wealthy children? *Journal of Political Economy 129*(3), 703–756.

Forrest, W. and C. Hay (2011). Life-course transitions, self-control and desistance from crime. *Criminology & Criminal Justice 11*(5), 487–513.

Fox, J., S. Kluesener, and M. Myrskyla (2019). Is a positive relationship between fertility and economic development emerging at the sub-national regional level? theoretical considerations and evidence from europe. *European Journal of Population 35*, 487–518.

Gale, L. S. and D. Shapley (1962). College admissions and the stability of marriage. *American Mathematical Monthly 69*, 9–15.

Giuntella, O., L. Rotunno, and L. Stella (2021). Trade shocks, fertility, and marital behavior. *SOEP papers on Multidisciplinary Panel Data Research, No. 1126*.

Goldin, C. (2021). *Career and Family: Women's Century-Long Journey Towards Equity*. Princeton University Press.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*(2), 254–277.

Hart, R. K. (2015). Earnings and first birth probability among Norwegian men and women 1995–2010. *Demographic Research 33*, 1067–1104.

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

Huttunen, K., J. Møen, and K. G. Salvanes (2011). How destructive is creative destruction? effects of job loss on job mobility, withdrawal and income. *Journal of the European Economic Association 9*(5), 840–870.

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993). Earnings losses of displaced workers. *American Economic Review 83*(4), 685–709.

Jalovaara, M., G. Neyer, G. Andersson, J. Dahlberg, L. Dommermuth, P. Fallesen, and T. Lappegård (2019). Education, gender, and cohort fertility in the Nordic countries. *European Journal of Population 35*(3), 563–586.

Jones, L. E., A. Schoonbroodt, and M. Tertilt (2011). Fertility theories: Can they explain the negative fertility-income relationship? In *Demography and the Economy*. University of Chicago Press.

Juhn, C. and K. McCue (2017). Specialization then and now: Marriage, children, and the gender earnings gap across cohorts. *Journal of Economic Perspectives 31*(1), 183–204.

Kalil, A., M. Mogstad, M. Rege, and M. E. Votruba (2016). Father presence and the intergenerational transmission of educational attainment. *Journal of Human Resources 51*(4), 869–899.

Kearney, M. S. and R. Wilson (2018). Male earnings, marriageable men, and nonmarital fertility: Evidence from the fracking boom. *The Review of Economics and Statistics 100*(4), 678–690.

Kitterød, R. H. and M. Rønsen (2013). Does parenthood imply less specialization than before? Tales from the Norwegian time use surveys 1980-2010. *Statistics Norway Discussion Paper No 757*.

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.

Kolk, M. (2019). The relationship between lifecourse accumulated income and childbearing of swedish men and women born 1940-1970. *Stockholm Research Reports in Demography 19*, 1–38.

Kravdal, Ø. (2002). The impact of individual and aggregate unemployment on fertility in Norway. *Demographic Research 6*, 263–294.

Kravdal, Ø. and R. R. Rindfuss (2008). Changing relationships between education and fertility: A study of women and men born 1940 to 1964. *American Sociological Review 73*(5), 854–873.

Kunze, A. (2020). The effect of children on male earnings and inequality. *Review of Economics of the Household 18*(3), 683–710.

Lappegård, T. and M. Rønsen (2013). Socioeconomic differences in multipartner fertility among Norwegian men. *Demography 50*(3), 1135–1153.

Lappegård, T., M. Rønsen, and K. Skrede (2011). Fatherhood and fertility. *Fathering: A Journal of Theory, Research & Practice about Men as Fathers 9*(1), 103–120.

Lindo, J. M. (2010). Are children really inferior goods? *Journal of Human Resources 45*(2).

Løken, K. V., M. Mogstad, and M. Wiswall (2012). What linear estimators miss: The effects of family income on child outcomes. *American Economic Journal: Applied Economics 4*(2), 1–35.

Lovenheim, M. F. and K. J. Mumford (2013). Do family wealth shocks affect fertility choices? evidence from the housing market. *Review of Economics and Statistics 95*(2), 464–475.

Low, C. (2022). Pricing the biological clock: The marriage market costs of aging to women. *Mimeo*.

Lundberg, S., R. A. Pollak, and J. Stearns (2016). Family inequality: Diverging patterns in marriage, cohabitation, and childbearing. *Journal of Economic Perspectives 30*(2), 79–102.

Massenkoff, M. N. and E. K. Rose (2022). Family formation and crime. *Mimeo*.

Rege, M., K. Telle, and M. Votruba (2007). Plant closure and marital dissolution. *Statistics Norway Discussion Paper No 514*.

Salvanes, K. G., B. Willage, and A. L. P. Willen (2021). The effect of labor market shocks across the life cycle. *CESifo Working Papers* (9491).

Sampson, R. J., J. H. Laub, and C. Wimer (2006). Does marriage reduce crime? a counterfactual approach to within-individual causal effects. *Criminology 44*(3).

Schaller, J. (2016). Booms, busts, and fertility testing the Becker model using gender-specific labor demand. *Journal of Human Resources 51*(1), 1–29.

Shenhav, N. (2021). Lowering standards to wed? Spouse quality, marriage, and labor market responses to the gender wage gap. *Review of Economics and Statistics 103*(2), 265–279.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.

MALE FERTILITY: FACTS, DISTRIBUTION AND DRIVERS OF INEQUALITY APPENDIX FOR ONLINE PUBLICATION

A.1 Alternative descriptive mechanisms

We explore alternative mechanisms that could explain the increasingly negative relationship of childlessness, and positive relationship of total fertility, with relative earnings rank: health, incarceration and missing data.

Health Outcomes A potential alternative mechanism linking relative earnings and fertility is health: those with lower earnings may also have poorer health, which may affect their ability to either attract a partner or physically to have a child. In addition, measures of health and physical attractiveness, such as BMI, have been shown to be correlated with men's value on the marriage market (Chiappori, Oreffice, and Quintana-Domeque 2012). To explore this possibility, we consider two measures of health: long-term disability, and health status at conscription for mandatory military service at age 18. Figure A.1 depicts the relationship between relative earnings rank and the average proportion of individuals registered as having a long-term disability at age 30.[26] Although there is a negative correlation between relative earnings rank and permanent disability, the overall rates of disability are substantially lower than the rates of childlessness seen in Figure 1. Equally important, there is no indication that young-age disability rates have increased over time among low earners and that such developments could explain their rising rates of childlessness.

Figure A.2 shows height and BMI at conscription, by earnings rank, for two representative cohorts. While height is correlated with relative earnings rank (an average gap of around 2cm between the lowest and highest earning men), BMI is not. However, the differences in height are so small as to make it unlikely that there is a health-driven relationship between earnings rank and fertility.

Incarceration Men at the lower end of the earnings distribution may be unable to have a family because they are incarcerated. Figure A.3 explores this possibility by plotting, for two representative cohorts, the fraction of men with a prison sentence by relative earnings rank, with incarceration observed at age 30.[27] Predictably, the rates are highest for the lowest earners, but on average extremely low and below one percent of the population. More importantly, there is no indication that the relationship has tilted over time with rising incarceration rates for low earners. Incarceration is unlikely to be a key mechanism behind the stylized facts on male fertility.

Data Quality We consider whether data quality, and in particular the notion of "missing dads", can plausibly explain higher rates of childlessness among low-income men. Specifically, it may be

[26] These data are first available from 1992, and we are not able to study disability at young ages for the oldest cohorts included in earlier figures.

[27] These data are not available for the oldest cohort included in earlier figures.

that these men are not present long enough in the lives of the female partners to be registered as fathers at the time the child is born. Figure A.4 shows the relationship between the fraction of birth records missing a father's name, and the woman's earnings rank - given that the fathers are missing, it is not possible to depict this relationship by the man's earnings rank. However, the rates of birth records with missing fathers are low overall, at 0.7% for the whole sample.[28] They are highest for the lowest earning women, being close to 3% in the bottom 5% and less than 1% in the top 5%. The rates have not changed substantially across the three representative birth cohorts depicted. Although this could explain some part of the male fertility patterns we see, it is unlikely to explain the very high rates of childlessness (over 70% in the most recent cohorts) that are present among the lowest earning men and the time pattern of rising rates of childlessness.

Figure A.1: Disability and earnings.

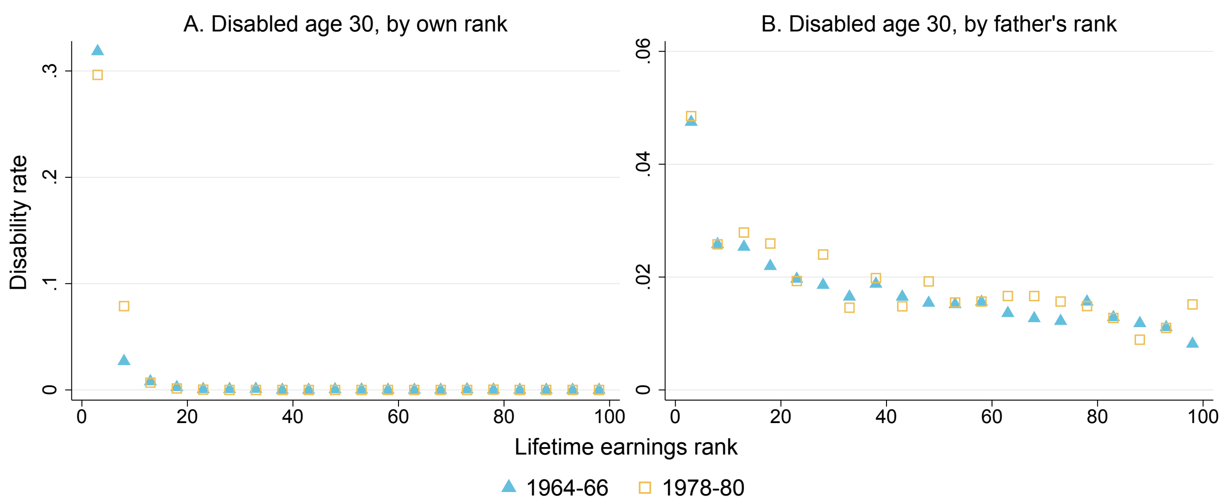

Notes: Each scatter point represents five percent of Norwegian men born between 1964-1966 and 1978-1980, respectively. Disability status is measured by receipt of a permanent disability pension at age 30. Observation count is 160 344. The average disability rate is 0.019.

[28]Some of these "missing dads" are in fact not missing, but are missing from the birth register because they do not have a Norwegian social security number.

Figure A.2: Earnings and other markers of health.

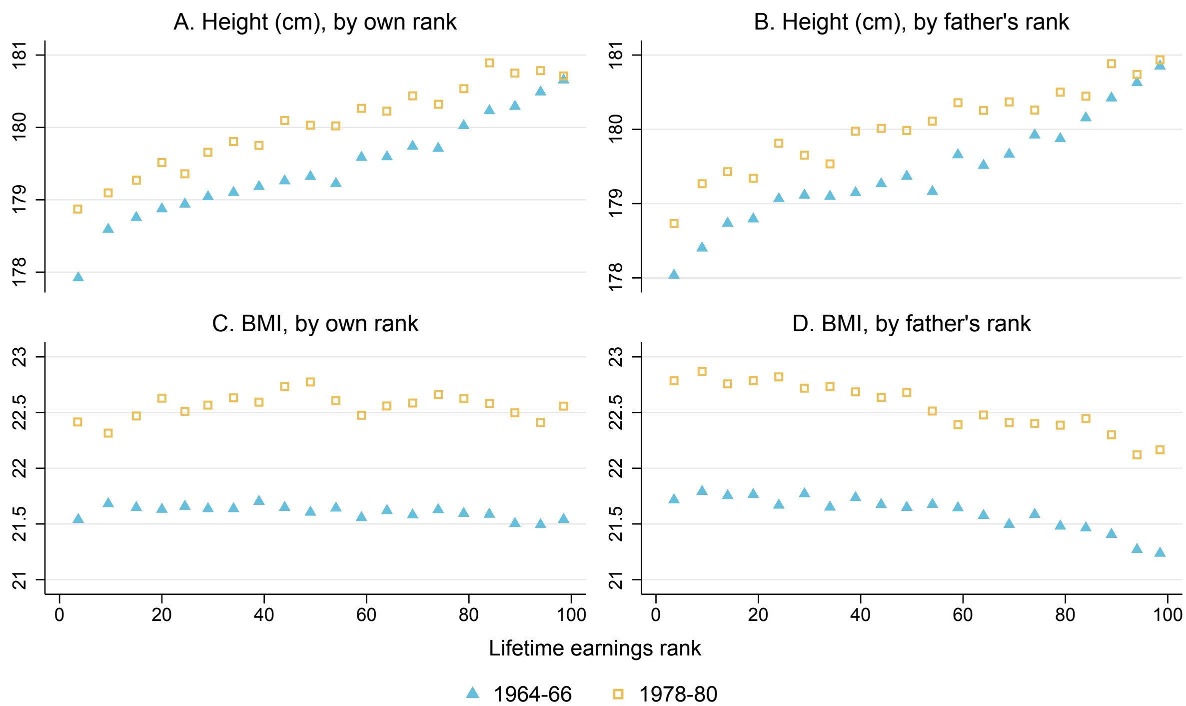

Notes: Each scatter point represents five percent of Norwegian men born between 1964-1966 and 1978-1980, respectively (data for the 1950-1952 cohorts are not available). Height and weight are measured at conscription for military service, typically at age 17 or 18. Observation count is 134 023. The average height is 179.6 cm, average BMI is 22.0, and mean obesity is 0.021.

Figure A.3: Incarceration and earnings.

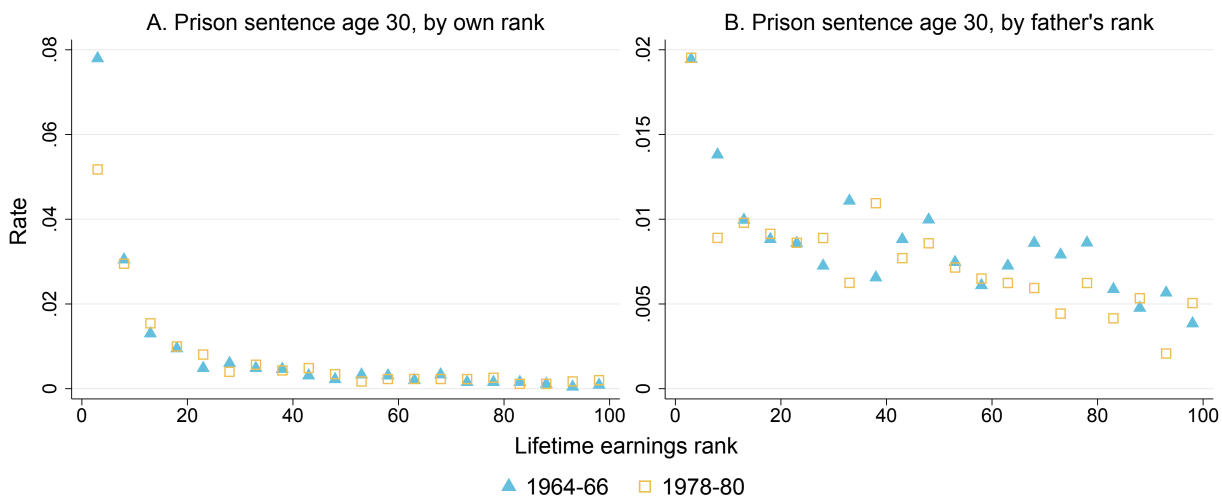

Notes: Each scatter point represents five percent of Norwegian men born between 1964-1966 and 1978-1980, respectively. Scatter points give the fraction of men charged with a crime and sentenced to unconditional imprisonment the year they turned 30. Observation count is 160 344. The average imprisonment rate is 0.0084.

Figure A.4: Missing birth records and mothers' earnings.

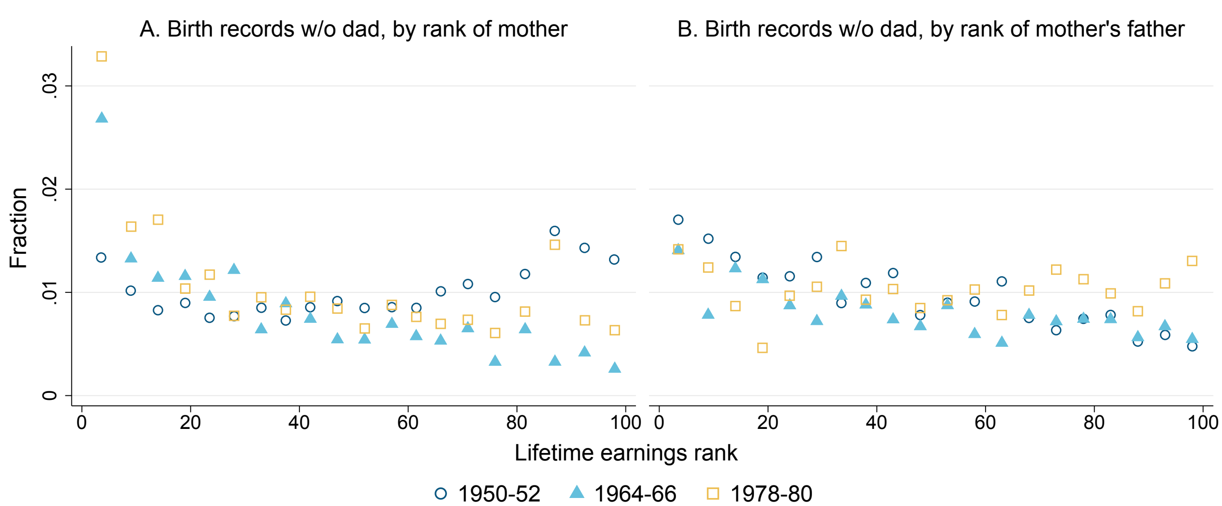

Notes: Each scatter point represents five percent of Norwegian women born between 1950-1952, 1964-1966 and 1978-1980, respectively, and with at least one birth between 1967 and 2020. Scatter points give the fraction of birth records with missing information on the child's father. Observation count is 471 759 children born to 206 443 women by age 40. Average rate is 0.0068 per birth record.

A.2 Additional figures

Figure A.5: Inequality in fertility over time, measured at age 50.

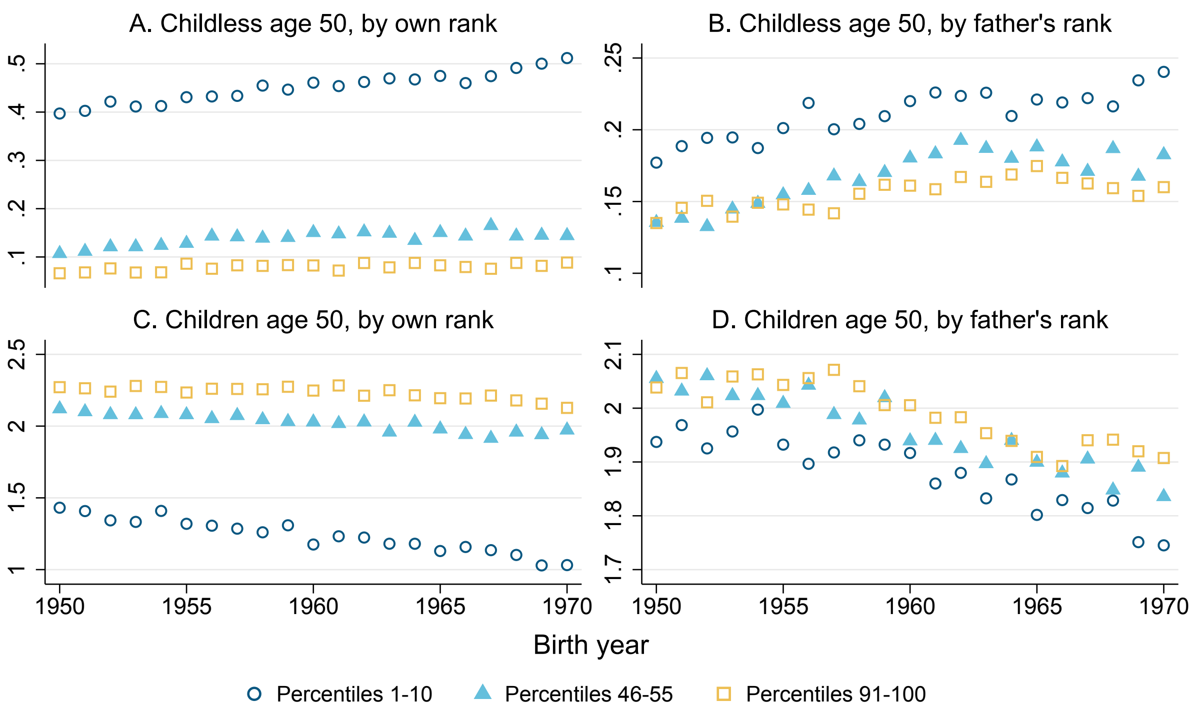

Notes: Scatter points represent ten percent of each cohort of Norwegian men born between 1950 and 1970.

Figure A.6: Partnership duration and out-of-wedlock births for the 1970-1980 male birth cohorts.

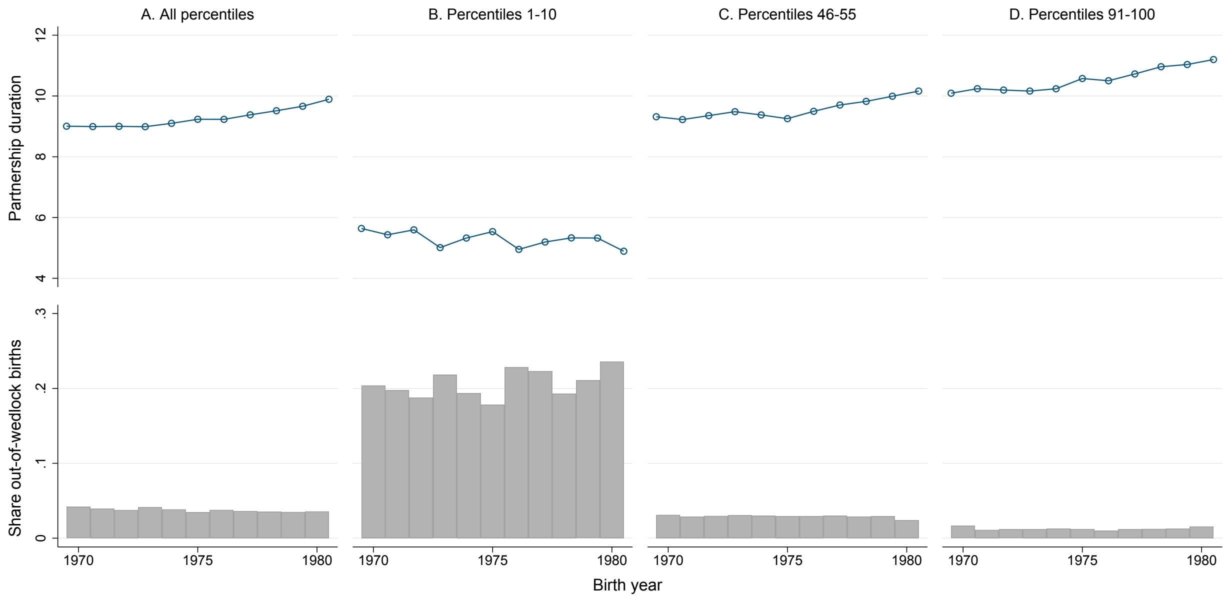

Notes: Computations of partnership duration and fraction of out-of-wedlock births draw on annual household records and partners as of end-of-year 1991-2020, restricted to age range 21-40. Panel A covers 244 776 unique partnerships and 468 085 child births; Panel B 14 618 unique partnerships and 21 948 child births; Panel C 25 994 unique partnerships and 50 034 child births; and Panel D 26 982 unique partnerships and 56 349 child births.

Figure A.7: Absolute earnings over time for women.

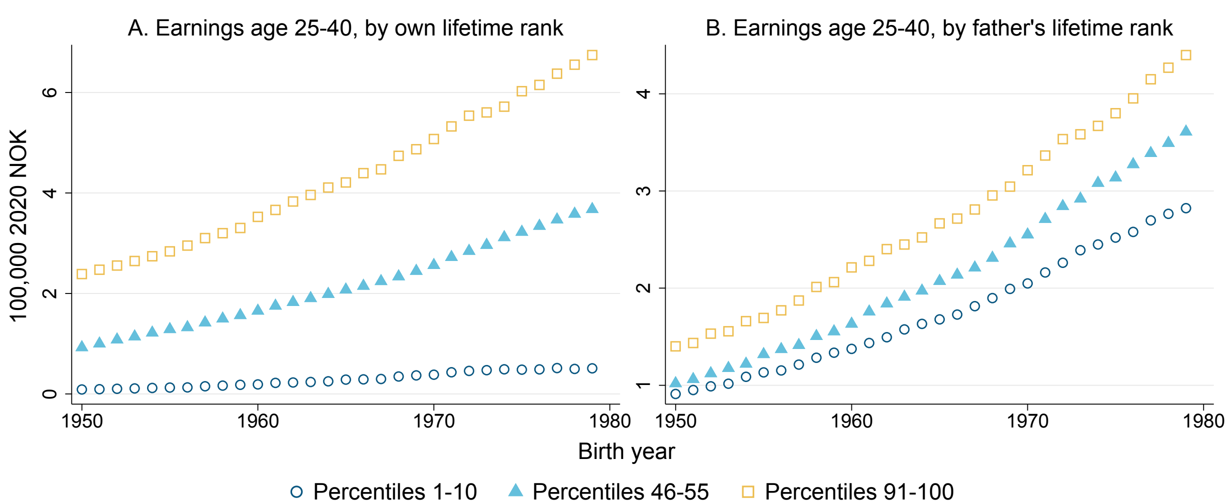

Notes: Scatter points represent ten percent of each cohort of Norwegian women born between 1950 and 1979. Depicted on the vertical axis are average annual earnings inflated to 2020 NOK, in units of 100 000. Observation count is 250 575.

Figure A.8: Comparing rates of single status using two sources of data.

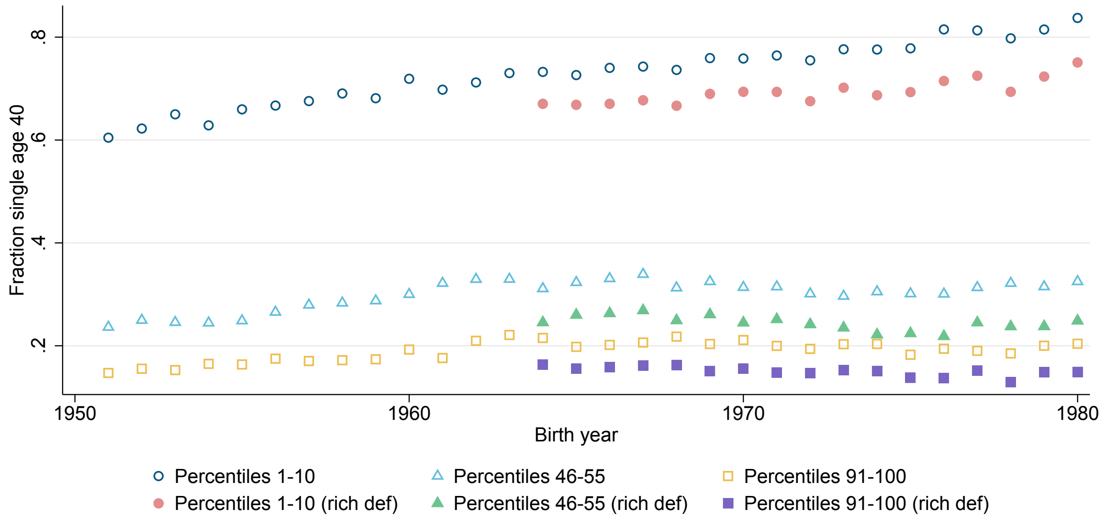

Notes: Measurement of single status draws on household records describing the household at the end of the calendar year. Single status as used throughout the paper is consistently measured as those not residing with a married partner nor cohabiting with a partner with whom the male has fathered a child. The broader ("rich") measure available since 2004 also excludes those cohabiting with a partner with a child fathered by someone else and those in childless partnerships.

Figure A.9: Fertility across the earnings distribution, women.

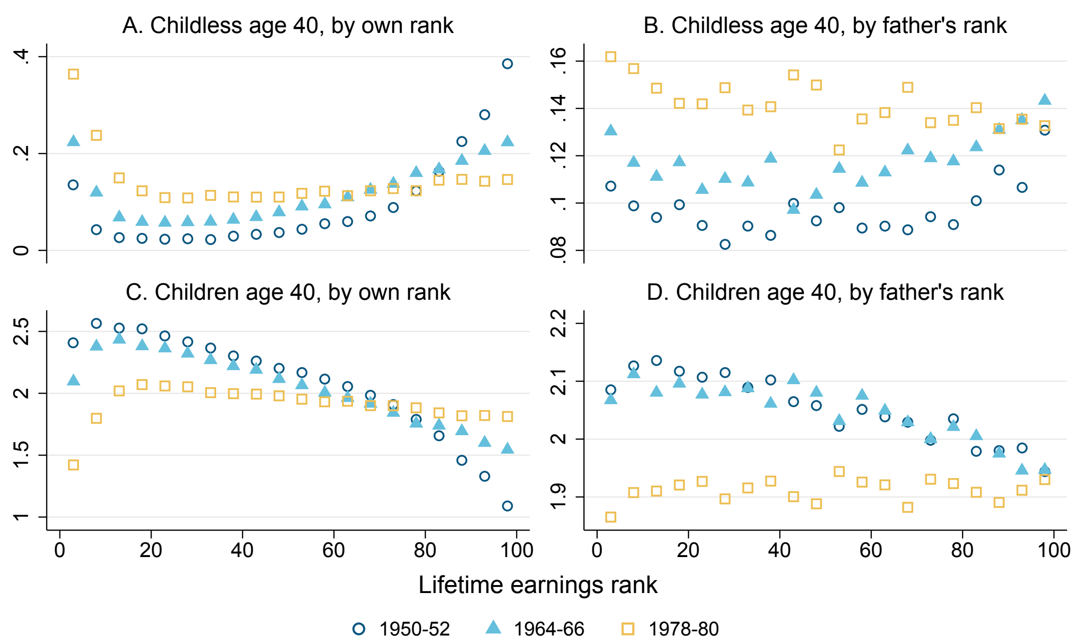

Notes: Each scatter point represents five percent of Norwegian women born between 1950-1952, 1964-1966, and 1978-1980, respectively. Observation count is 233 725.

Figure A.10: Father's rank by own earnings rank and multipartner status.

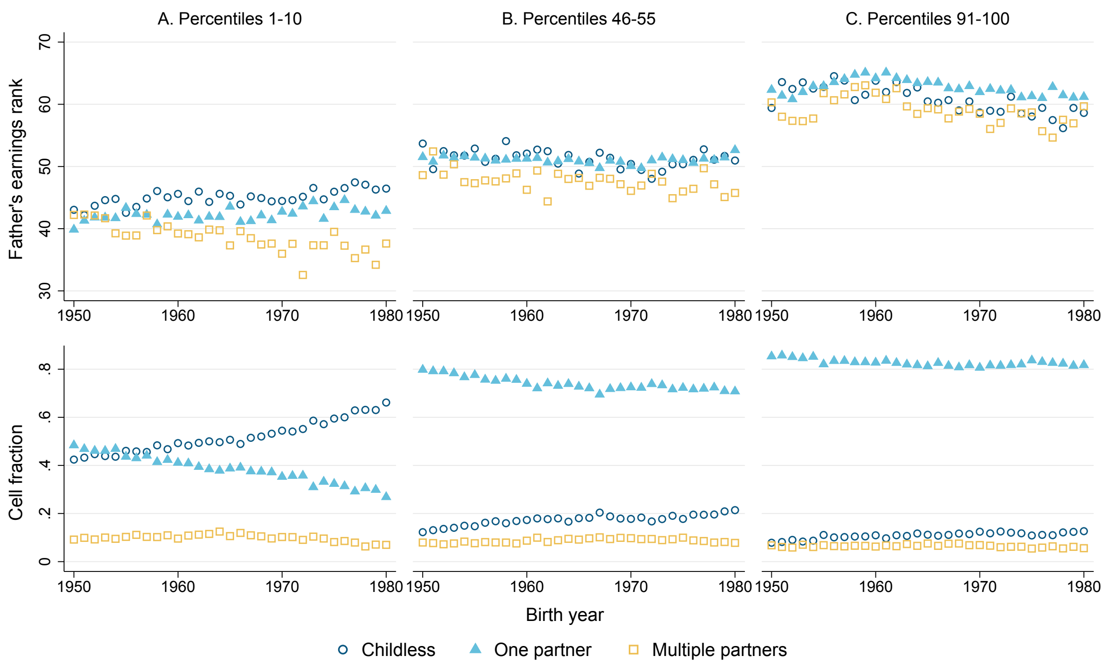

Notes: Scatter points in top panels show average of father's lifetime earnings rank for groups of men separated by own lifetime earnings rank (bottom, middle, and top deciles), birth year, and multipartner status by age 40. Bottom panels show the fractions of multipartner status within earnings decile and birth year.

Figure A.11: Extensive and intensive margin partnership for percentiles of the earnings distribution, across cohort.

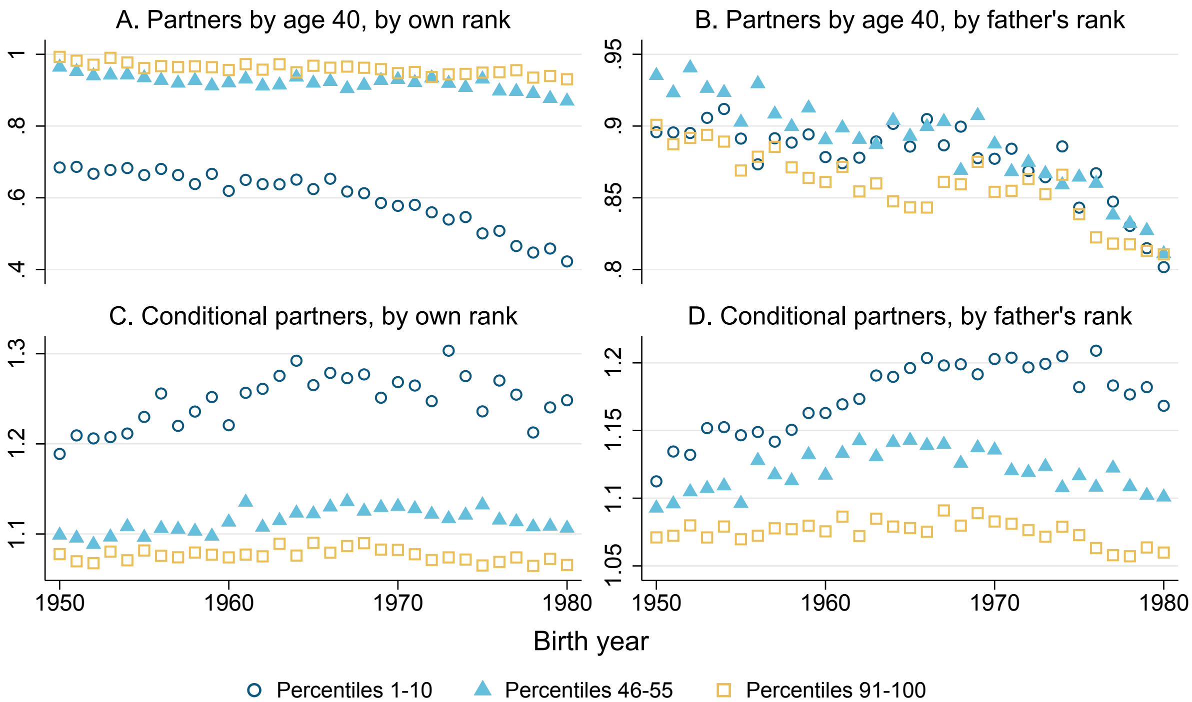

Notes: Population in Panels C and D is restricted to those with at least one child by age 40. See also notes to Figure 4.

Figure A.12: Sibling mean comparisons before and after firm bankruptcies, comparing singles and couples at t-2.

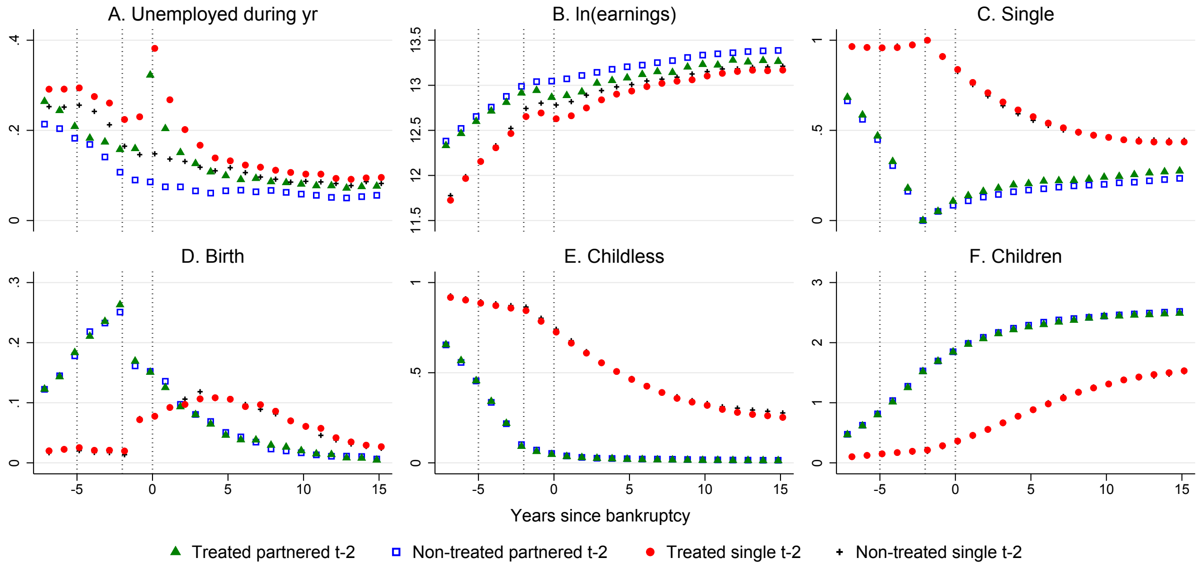

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). The sample is split at t-2 into single men and those with partners.

Figure A.13: Impacts of firm bankruptcies, comparing singles and couples at t-2.

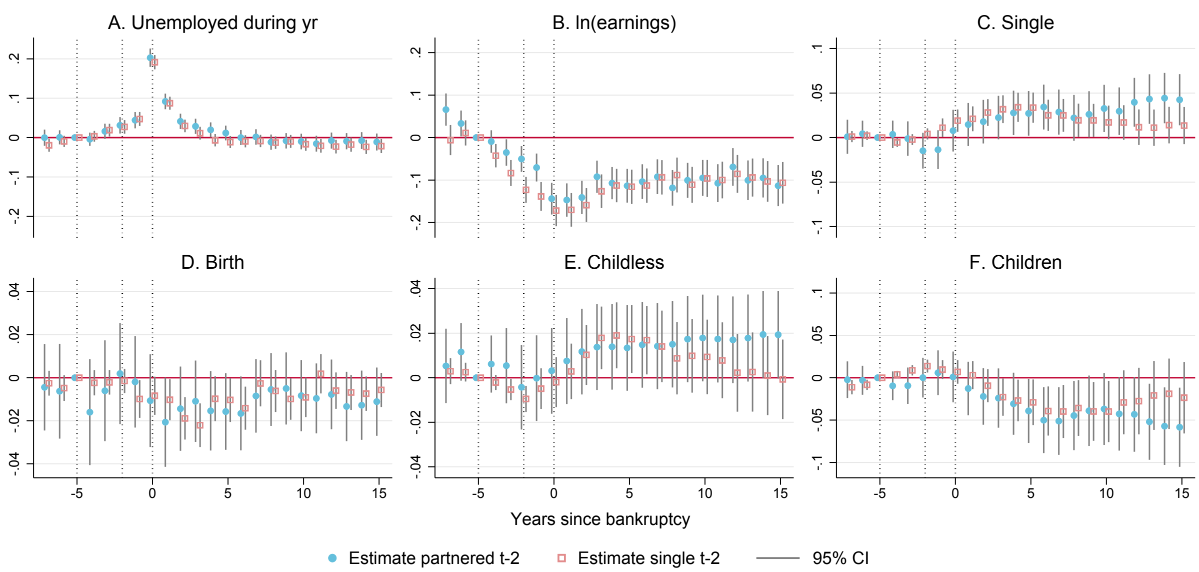

Notes: Regression model is augmented with a full set of interaction terms between an indicator for single status at time -2 and the treatment and time indicators of the baseline model.

Figure A.14: Sibling mean comparisons before and after firm bankruptcies, sampling at t-3.

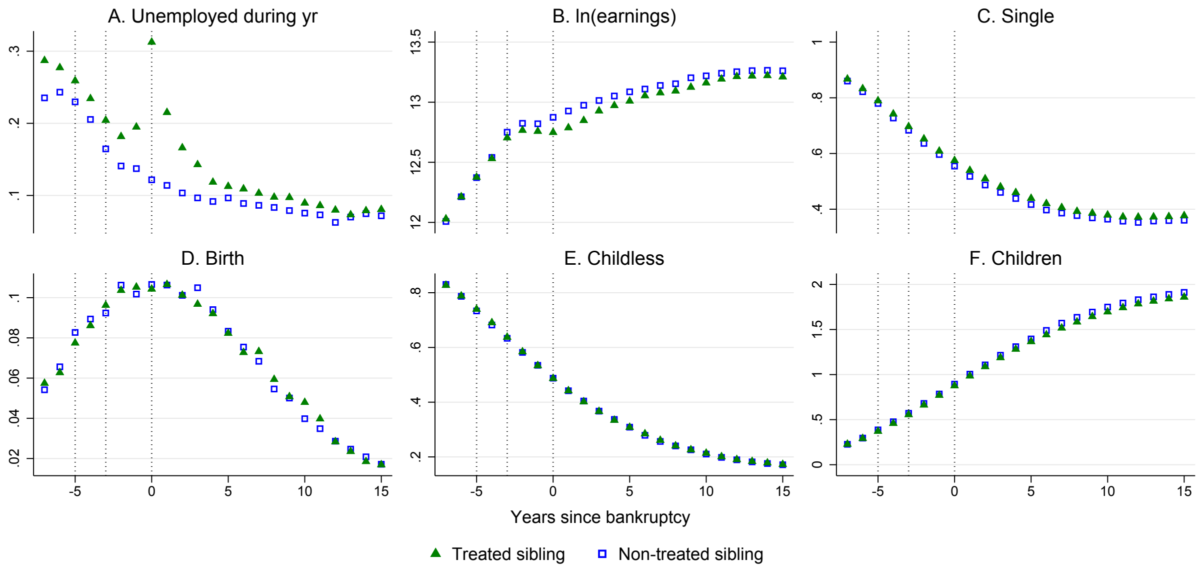

Notes: Vertical lines indicate year of observed November job (year -3), year of event (year 0), and reference year (-5). See text and notes to A.15 for a description of samples.

Figure A.15: Effects of firm bankruptcies, sampling at t-3.

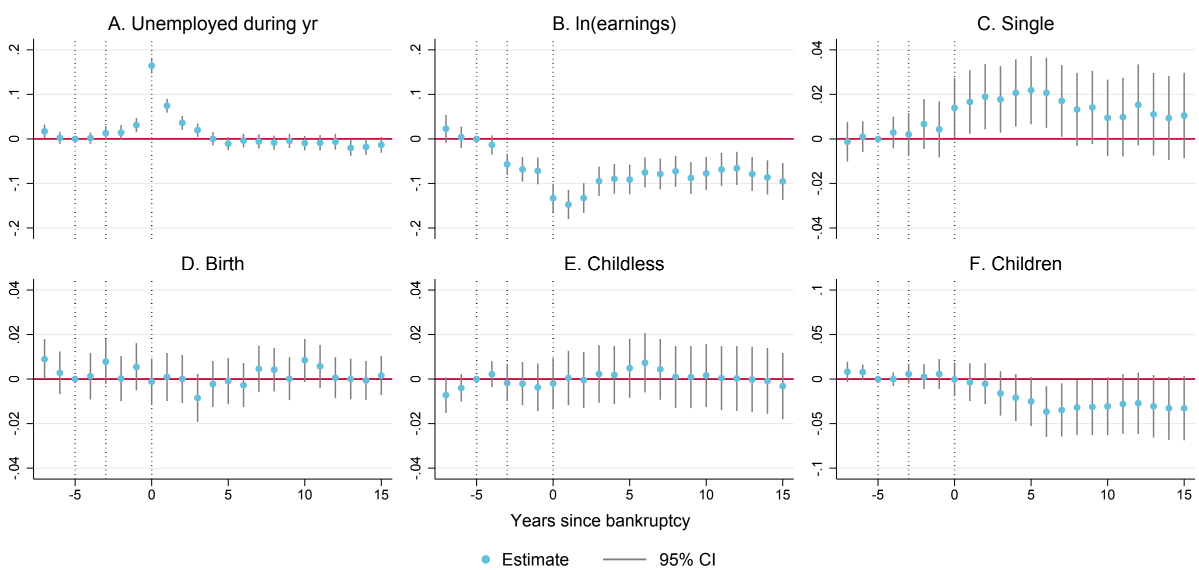

Notes: Vertical lines indicate year of observed November job (year -3), year of event (year 0), and reference year (-5). Scatter points show the estimates of β_t from the estimating equation. Sample of treated siblings consists of Norwegian-born men who in year -3 worked in a firm that filed for bankruptcy three years later and were age 25-35 in the year of the event, while non-treated siblings in year -3 held a job in a firm that did not file for bankruptcy during the observation period. Samples are restricted to families with both treated and non-treated siblings. Observation counts are 213 202 in the treatment group and 251 919 in the control group.

Figure A.16: Fadlon-Nielsen alternative control sampling strategy.

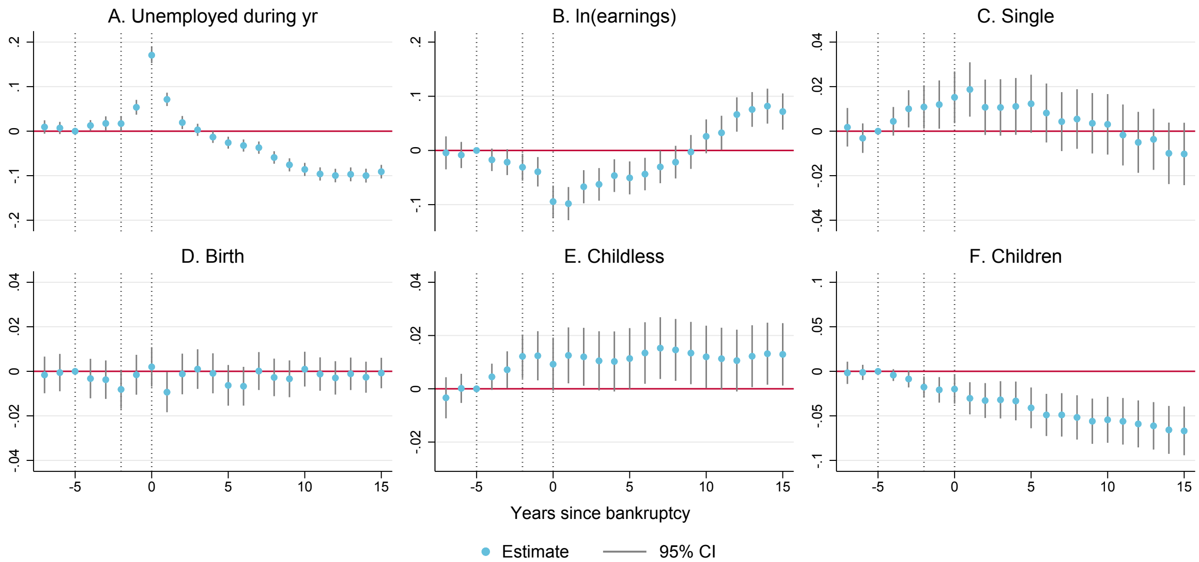

Notes: Estimates using an alternative control sample of men who experience a bankruptcy in future years (2006-2019), but are used as a control sample in the pre years when they have not yet experienced that bankruptcy. The control sample is matched to a treatment sample of men who were age 25 to 35 and experienced a bankruptcy between 1995 and 2005. Based on Fadlon and Nielsen (2019) approach. Sample sizes are 391,079 in treatment and 326,555 in control group respectively.

Figure A.17: Effects of firm bankruptcies, accounting for family-by-year fixed effects.

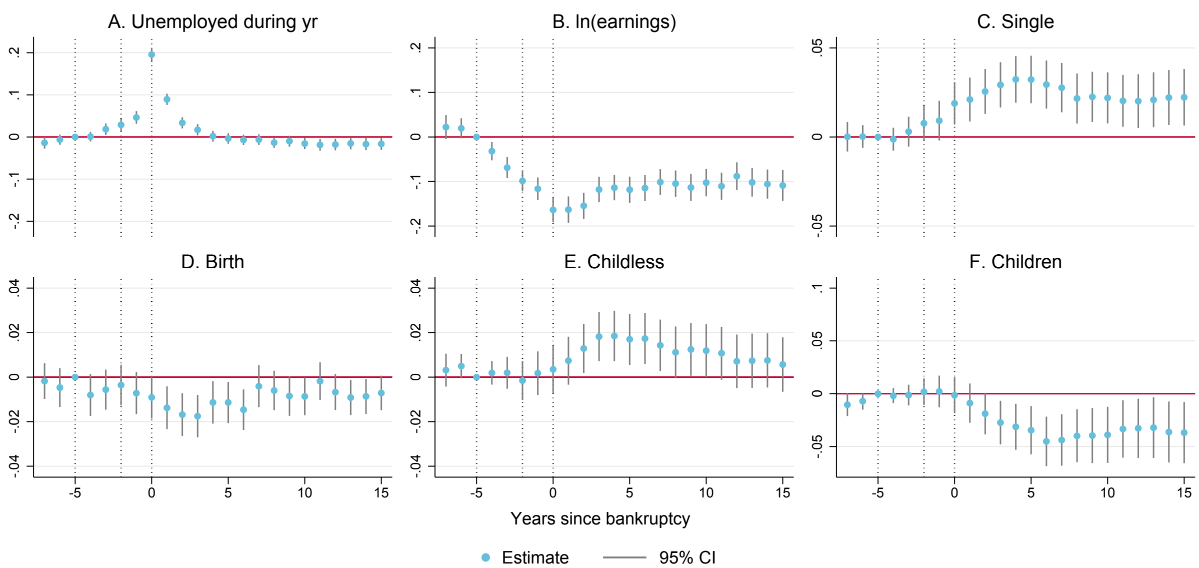

Notes: Regression model is augmented with family-by-year fixed effects. See also note to Figure 8.

Figure A.18: Sibling mean comparisons before and after establishment shutdowns.

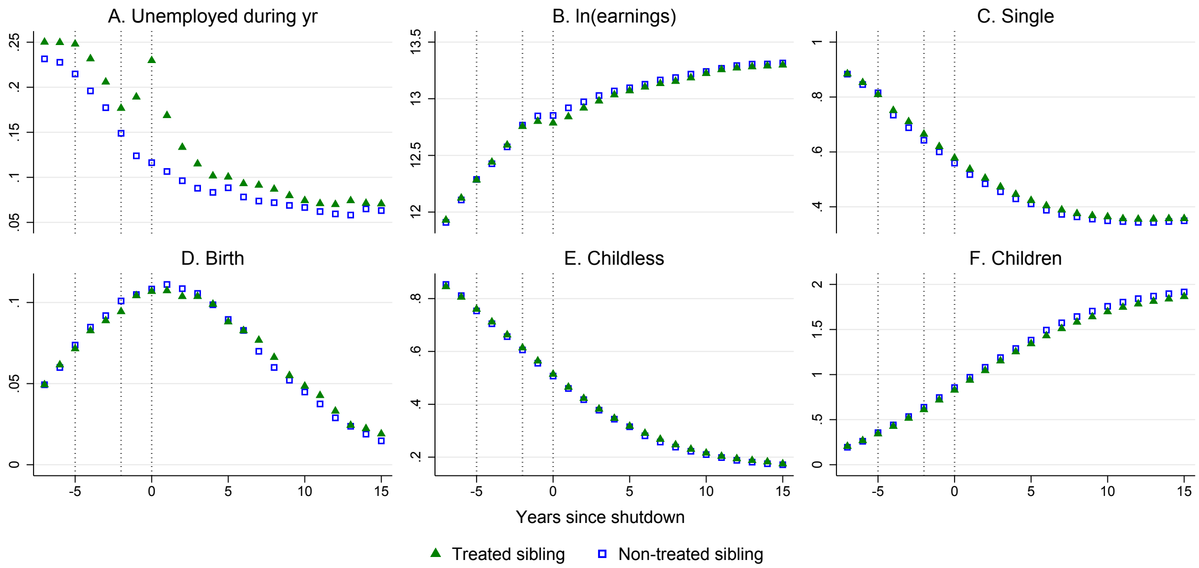

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). See text and notes to Figure A.19 for a description of samples.

Figure A.19: Effects of establishment shutdowns.

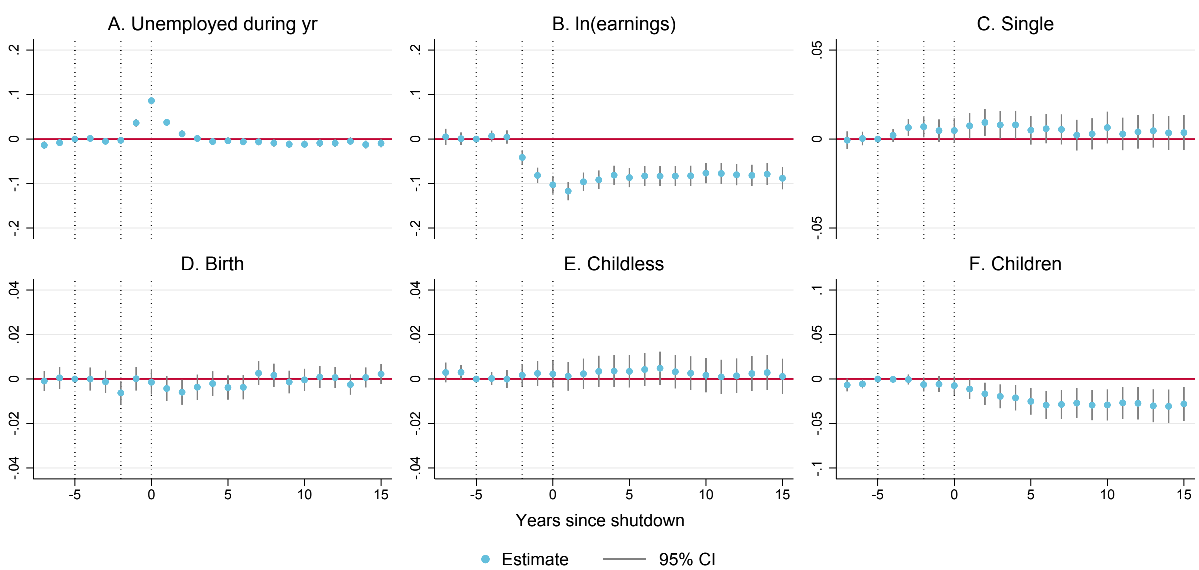

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Scatter points show the estimates of β_t from the estimating equation. Sample of treated siblings consists of Norwegian-born men who in year -2 worked at an establishment that shut down two years later (between 1995 and 2015) and were age 25-35 in the year of the event, while non-treated siblings in year -2 held a job in an establishment that did not shut down during the observation period. Samples are restricted to families with both treated and non-treated siblings. Observation counts are 797 144 in the treatment group and 940 403 in the control group.

Figure A.20: Removing any alternative bankruptcy events.

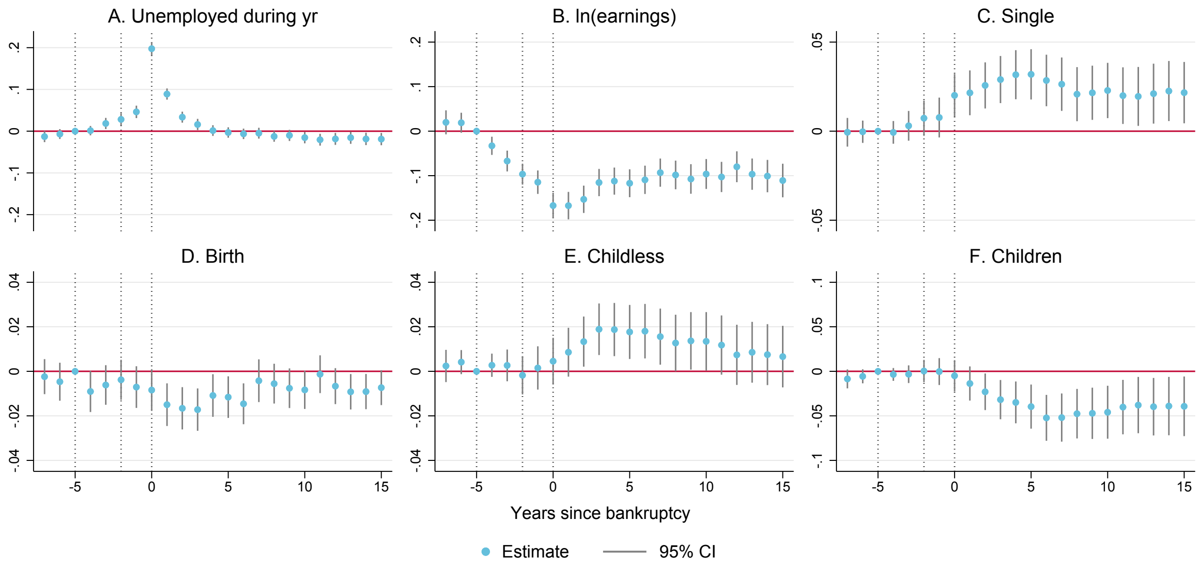

Notes: Regression samples exclude individuals who experience bankruptcy in other years than yr 0, so that the treatment sample is restricted to individuals who experience only one bankruptcy and the control sample to individuals who do not experience a bankruptcy during the observation window. Observation counts are 263 349 in the treatment group and 315 269 in the control group.

Figure A.21: Removing large bankruptcies.

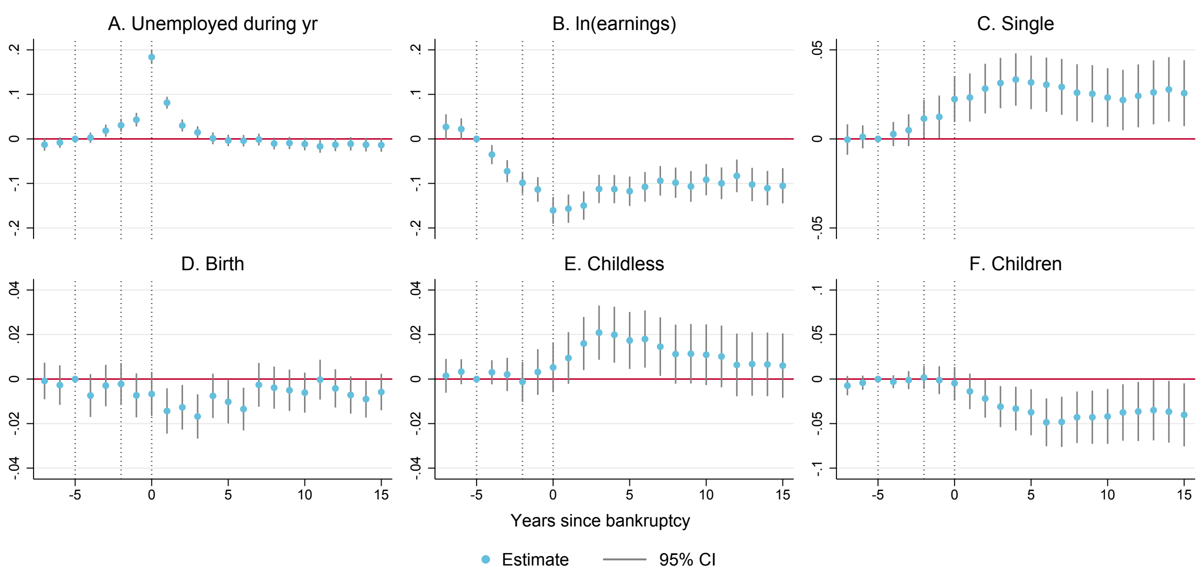

Notes: Regression samples exclude event sequences where the individual in year -2 resided in a municipality where at least 1.5 percent of the workforce worked in a firm that filed for bankruptcy two years later. Observation counts are 237 915 in the treatment group and 282 042 in the control group.

Figure A.22: Removing small bankruptcies.

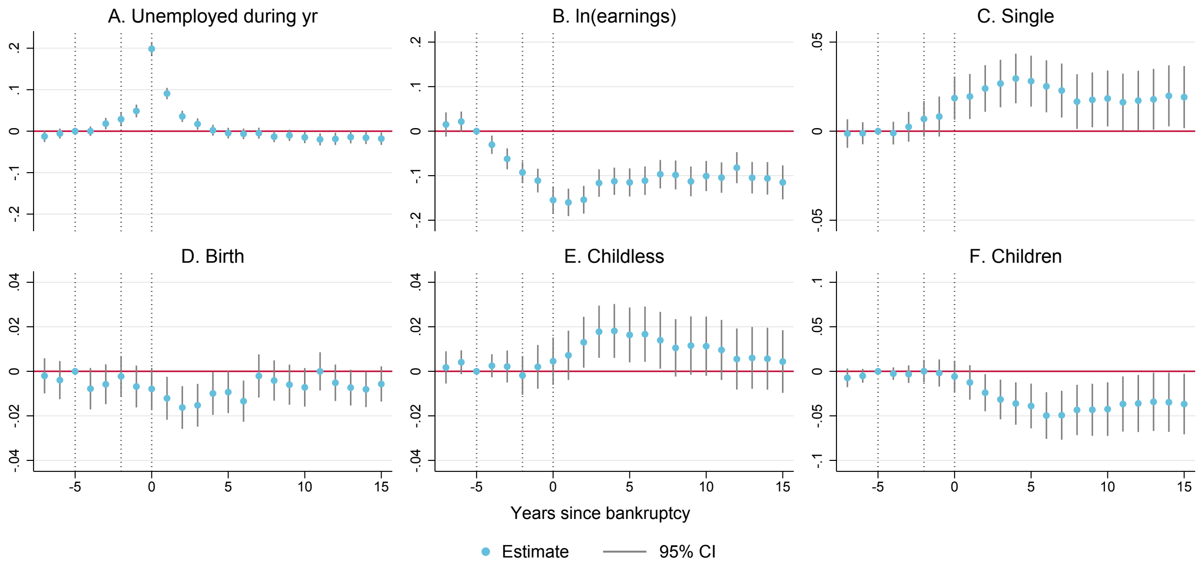

Notes: Regression samples exclude individuals in the treatment group who in year -2 worked in a firm with fewer than six employees. Observation counts are 250 281 in the treatment group and 317 579 in the control group.

Figure A.23: Sibling mean comparisons before and after bankruptcies, balanced sample.

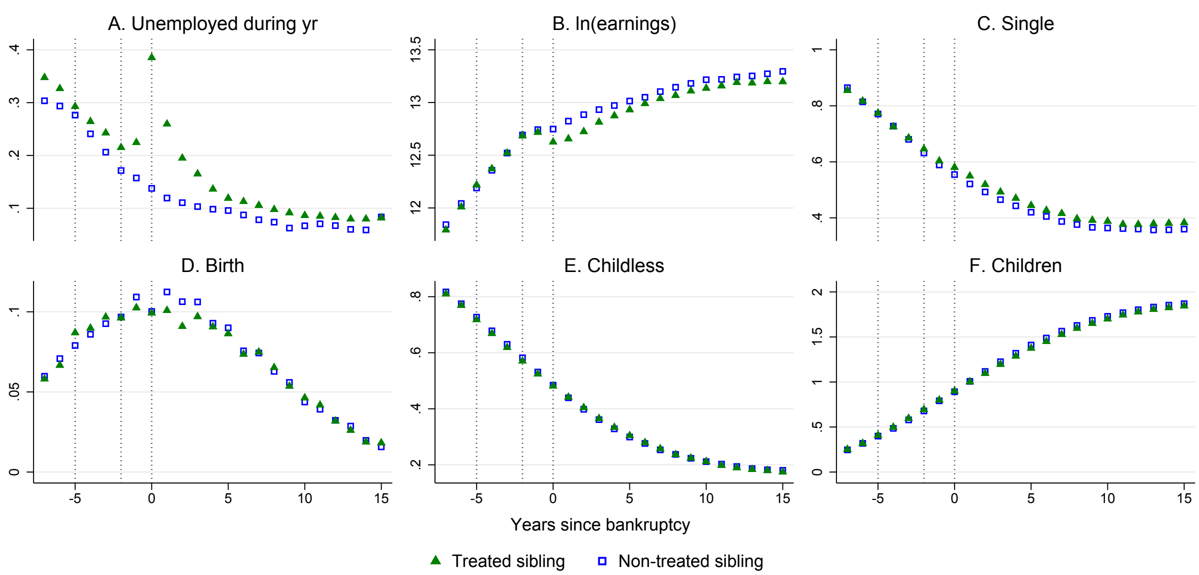

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Sample of treated siblings consists of Norwegian-born men who in year -2 worked at an establishment that shut down two years later (between 1997 and 2005) and were age 25-35 in the year of the event, while non-treated siblings in year -2 held a job in an establishment that did not shut down during the observation period. Samples are restricted to families with both treated and non-treated siblings. Observation counts are 132 186 in the treatment group and 166 265 in the control group.

Figure A.24: Effects of firm bankruptcies, balanced sample.

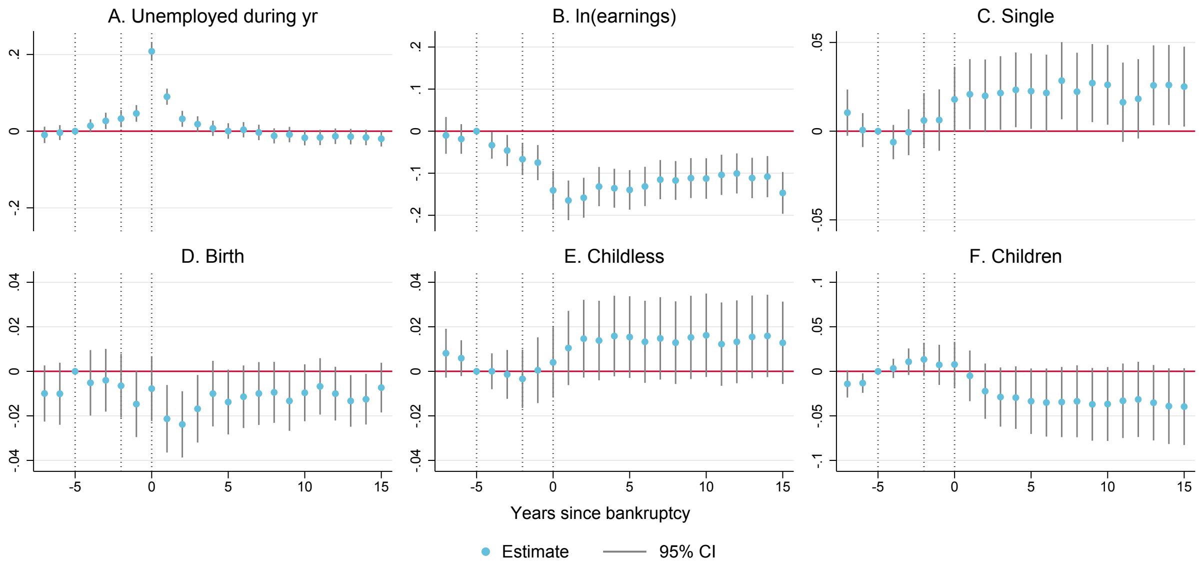

Notes: Regression model is estimated on restricted balanced sample where the 23-year time sequence falls within the data window 1990-2020, with bankruptcies in 1997-2005. Observation counts are 132 186 in the treatment group and 166 265 in the control group.

Figure A.25: Comparing our main estimates to estimates from a stacked regression approach.

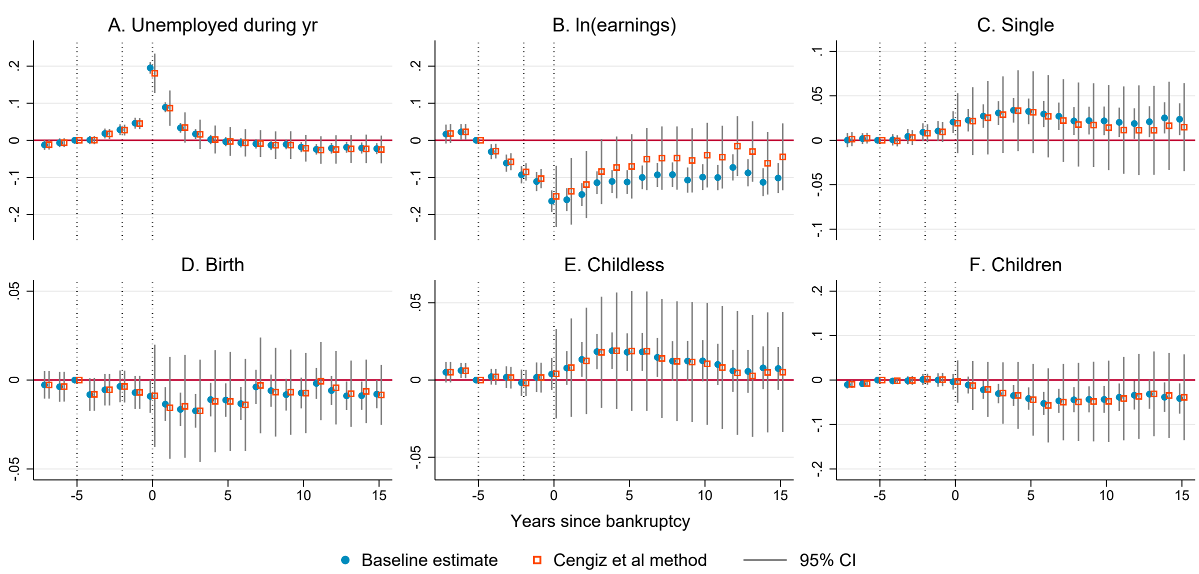

Notes: Baseline estimates replicate those in Figures 7 and 9. The Cengiz et al approach draws on a stacked event-by-event analysis, where each point estimate is based on 21 separate regressions omitting any observations where already treated individuals may influence estimation of calendar year effects.

Figure A.26: Average education and unemployment over time

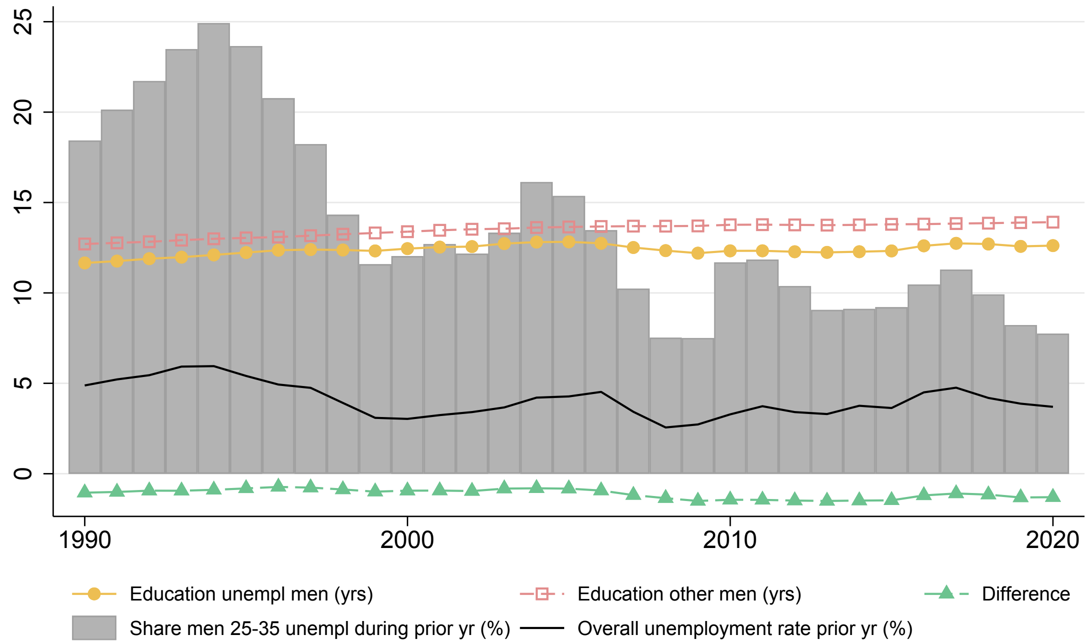

Notes: Vertical bars show the fraction of men aged 25-35 with registered unemployment during the prior year; and solid line shows the lagged annual unemployment rate collected from https://data.oecd.org/unemp/unemployment-rate.htm. Scatter points depict mean years of schooling for those with and without registered unemployment, as well as the difference in attainment. Observation count is 9 162 238.

Figure A.27: The relationship between education and fertility over time, sensitivity to controls

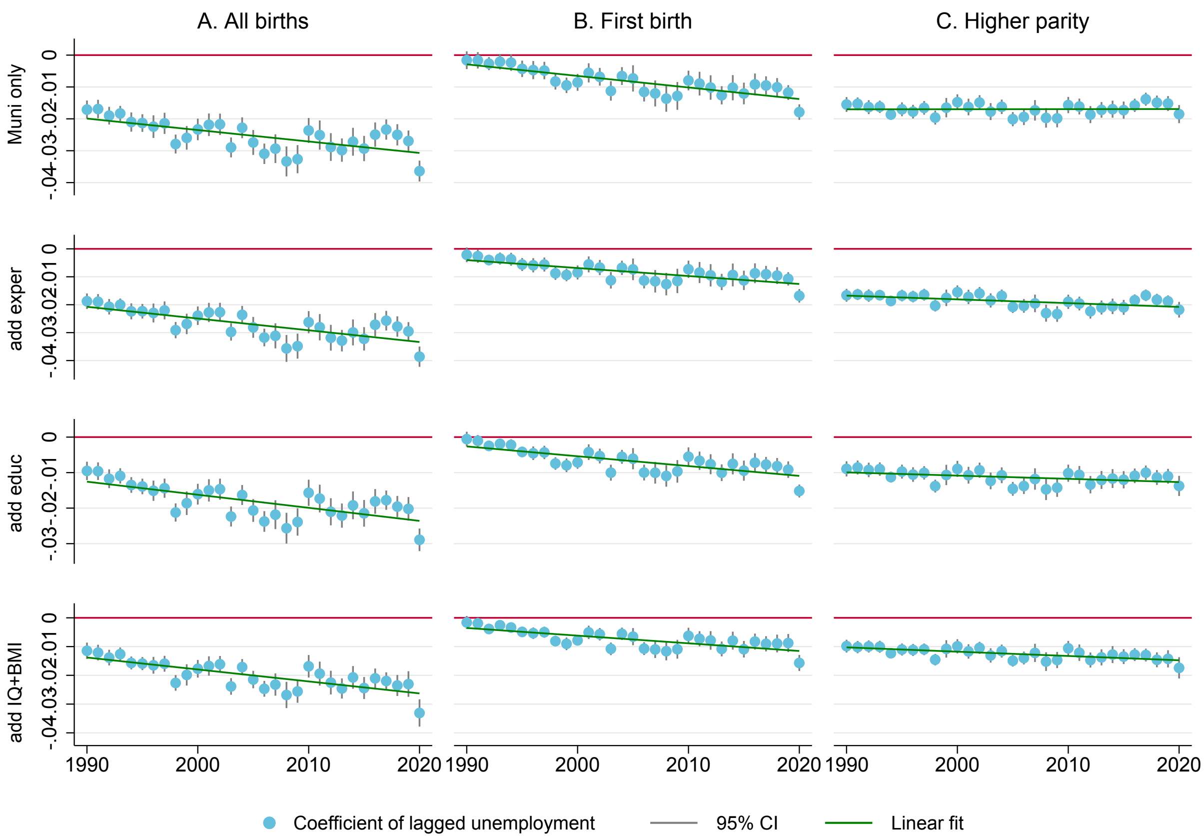

Notes: Scatter points show the estimated coefficient of individual unemployment status from a regression of birth on registered unemployment the prior year. All regressions control for year of observation. Regression in top row includes 428 municipality fixed effects; the second row adds polynomial of years of experience; the third row adds educational attainment to the specification of the second row; and the final row adds IQ and a polynomial of BMI to the model of the third row. Standard errors are clustered within municipality. Sample consists of men age 25-35, sample period is 1990-2020. Observation count is 9 162 238 (7 966 882 in bottom panels because of missing IQ or BMI data). Mean birth rate is 0.098 and mean registered unemployment is 0.141.