Male Fertility: Facts, Distribution and Drivers of Inequality

Bernt Bratsberg* Andreas Kotsadam† Selma Walther‡

October 25, 2021

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 show that the gap in male childlessness between low and high earners has widened by almost 20 percentage points over the last thirty years. Using firm bankruptcies, we provide evidence that men experiencing negative labor market shocks are less likely to become fathers and be partnered, and that these effects are persistent up to 15 years. We conclude by documenting that men's fertility penalty to job loss has increased markedly over the last three decades.

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 Jesper Bagger, Amalavoyal Chari, Monica Costa Dias, Elisabetta De Cao, Paola Giuliano, Chris Hansman, Hamish Low, Vegard Skirbekk, and seminar participants at Oxford, Maastricht, Sussex, Norwegian Institute of Public Health, the SEHO 2021 Meeting in Boston, 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

There have been a number of societal developments in recent decades in high income countries that have affected labor market opportunities and family outcomes, with potentially unequal impacts in the population. Marriage rates have declined, and divorce rates have increased. There has also been a growth in out-of-wedlock births and single mothers, and more generally a changing lifecycle pattern of women, with a reduction in the gender wage gap and increasing economic independence (Lundberg, Pollak, and Stearns 2016, Goldin 2004). The traditional role of the male breadwinner has been eroded, and women are less reliant on men to support and raise their families, partly due to improvements in women's labor force participation and pay.

Another key challenge faced by workers is a widening gap in the returns to skilled and unskilled labor, which has been particularly evident among men (Hornstein, Krusell, and Violante 2005). There have also been prominent trends by gender: while the returns to women's work have steadily grown over time, the returns to men's time, and particularly unskilled men's time, have stagnated and even declined in recent years. Recently, there is growing evidence for a group of "left behind" men with poor labor market, education and health outcomes (see, e.g. Autor and Wasserman 2013 and Binder and Bound 2019).

These changes are likely to have affected family formation and fertility. While much is already known about female fertility, and particularly on the trade-off between career and family in various contexts (see, e.g., Kleven, Landais, and Soegaard 2019, Adda, Dustmann, and Stevens 2017, Bhalotra, Venkataramani, and Walther 2021), relatively little is known about male fertility. In this paper, we bring our focus to male fertility. We use Norwegian registry data, which provides data on all births to the entire population of Norwegian men and women since 1967. The data is comprehensive, with only 0.7% of births to native women missing a father's name. This allows us to directly analyze male fertility using data on men, rather than indirectly using data on women's births, as the extant literature has done. As part of this dataset we also have access to a rich set of labor market outcomes and other family outcomes.

We begin by documenting two new stylized facts. First, we show that childlessness is highest among the poorest men, as captured by their within-cohort earnings rank. In particular, while childlessness rates are 72% among the bottom 5% of the earnings distribution in the most recent cohorts, they are only 11% among the top 5% richest men. Second, we document that this inequality in fertility has widened over time. Overall rates of childlessness have increased for all men over time, but they have increased more for the lowest earning men. These developments are not explained by changes in fertility delay. There has been a compression of the fertility distribution, with fewer men experiencing a larger share of the population's new births.

There are several possible mechanisms that may explain these patterns. We show descriptively that there is evidence for the importance of economic reasons, with patterns of men's relative earnings mimicking those of male childlessness. Lower earning men are also more likely to be single, indicating a role for the marriage market. There is less evidence to support the importance of health reasons: we do not see similar patterns in disability, height or BMI. There is also little evidence for data quality issues, with only 0.7% of birth records having "missing dads".

Next, we use a robust empirical strategy to document this relationship between male earnings and fertility in a more causal way. We use firm bankruptcies as a shock to male employment and earnings (Bratsberg, Raaum, and Røed 2018) in an event study approach that conditions on individual and cohort*year fixed effects, follows individuals for seven years before and 15 years after the bankruptcy event and includes same-sex siblings as a comparison group.[1] Strikingly, we see remarkable persistence in the negative impacts of job loss on labor market and family life. As well as a heightened unemployment risk and reduced long-run earnings, men who experience job loss have fewer children overall, are more likely to still be childless, and are less likely to be partnered, 15 years after experiencing a firm bankruptcy. Interestingly, the probability of experiencing the birth of 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 never compensated for in later life. A back-of-the-envelope calculation indicates that between 41%-46% of the factual patterns of childlessness and total fertility can be explained by a causal earnings-fertility relationship.

We confirm that our 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 unchanged by 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.

Our results on firm bankruptcies show that men experiencing earnings losses are less likely to become fathers, but do not speak to how the relationship between male labor market prospects and fertility has changed over time. To dig deeper into the changing nature of family prospects among low earners, we estimate the descriptive correlation between job

[1] 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.

loss in the previous year, as proxied by the individual claiming unemployment benefits, and having a child the following year, conditioning on a wide set of covariates, for each calendar year between 1990-2019. 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 than other men, the crucial finding is that the magnitude of this effect has become larger over time. We also show that this pattern is not explained by changes in the composition of the unemployed over time. This provides further evidence for the notion of "left behind" men: in recent years, men with poor labor market outcomes are facing stronger penalties in family outcomes, and specifically fertility.

We contribute most closely to the budding literature on the economic and family outcomes of men. Adding to evidence of "left behind" men, Autor, Dorn, and Hanson (2019) use a shift-share instrument in the U.S. deriving from Chinese import shocks to study the impact of reductions in males' relative earnings on a selection of male outcomes, and find that young adult men are particularly negatively affected by trade shocks. They also find increases in single motherhood and male premature mortality, and a reduction in male marriage and fertility, but using fertility data on women.[2] 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. They find that income growth promotes both marital and non-marital childbearing, but identify this through data on the fertility outcomes of women.

We take an important step forward in this literature by showing that fertility is another, crucial dimension along which low earning men are being "left behind". A limitation of the existing literature is that it primarily uses data on women: this does not speak to which men are having families, and how this may have changed over time, a key question that we tackle. By using direct data on male fertility for an entire population over several decades, we are able to document inequality in fertility between men, show that the population's new births are occurring to a shrinking fraction of the male population over time, and that poor labor market prospects play a key role in these changes.

We also contribute to the established literature on the determinants of fertility, and in particular how fertility responds to changes in income. 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

[2] 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.

correlate with the probability of finding a partner. Hence, it is likely that job loss affects partnering and by focusing on couples the identified effects are limited to only a selected subset. 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. Black, Kolesnikova, Sanders, and Taylor (2013) take a slightly different approach, focusing 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 the demand for coal in the Appalachian coal-mining region in the 1970s led to more births. While our focus here is distinct from this literature, zooming in on male fertility, how it is distributed across the population and how this inequality has changed over time, we see clear parallels in our empirical approach that uses bankruptcies as an exogenous shock to male earnings. Similar to Lindo (2010), we also include individual fixed effects in our estimation models, and as in Huttunen and Kellokumpu (2016), we take care to use an appropriate comparison group to minimise any bias arising from selection into firms.

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 link the descriptive and causal evidence of the paper, by providing a back of the envelope calculation of the share of the descriptive relationship between labor market outcomes and fertility that is likely to be causal, and showing that the relationship between labor market outcomes and fertility has changed over time. Section 7 concludes.

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 lead 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 has increased and that the Norwegian universal childcare allowance, which is universally given to all parents, has fallen in real terms. These factors may affect fertility negatively.

In terms of labor market policies, there have not been any dramatic changes and 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.

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.

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 1951. 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.

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 16 through 50 for individuals born between 1951 and 1969, and through age 40 for those born 1951-1979. 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 sample* using 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 and event study control samples to families represented with same-sex siblings in both samples.

Finally, in the *Stacked cross-sectional samples* used in Section 6.2, we pool cross-sectional population data for the period 1990-2019 and study changes across time in the correlation between individual unemployment status and 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.[3] In our population sample of men, about 7.5 percent were registered as unemployed in a given year.

[3]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.

Table 1: Descriptive statistics, male samples

	Population (1)	Stacked cross-sectional (2)	Event study (treated siblings) (3)	Event study (control siblings) (4)
Childless	0.213 [0.410]	0.551 [0.497]	0.440 [0.496]	0.440 [0.496]
Birth	0.041 [0.199]	0.098 [0.297]	0.075 [0.263]	0.076 [0.265]
First birth	0.009 [0.095]	0.046 [0.208]	0.032 [0.177]	0.032 [0.176]
Children	1.747 [1.209]	0.762 [0.986]	1.079 [1.163]	1.101 [1.183]
Single	0.327 [0.469]	0.611 [0.487]	0.554 [0.497]	0.531 [0.499]
Unemployed (during year)	0.075 [0.263]	0.143 [0.350]	0.176 [0.381]	0.127 [0.333]
Lifetime earnings rank	50.7 [28.8]	50.1 [26.0]	45.7 [22.4]	50.7 [24.3]
Other characteristics				
Age	40.0 [.]	30.0 [3.2]	32.8 [7.0]	33.0 [7.1]
Education (years)	13.3 [2.6]	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 [3.0]	22.3 [3.0]	22.2 [3.0]
Father's lifetime earnings rank	50.5 [28.8]	51.7 [23.1]	48.0 [22.1]	48.0 [22.1]
Birth year	1964.7 [8.2]	1973.7 [9.2]	1974.1 [5.9]	1973.7 [7.0]
Observation year	2004.7 [8.2]	2003.8 [8.6]	2006.8 [7.2]	2006.8 [7.4]
Age range	40	25-35	18-50	18-50
Observation period	1991-2019	1990-2019	1988-2019	1988-2019
Observations	816 535	8 881 215	262 865	311 608
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.

3 Stylized Facts

We begin by documenting patterns of fertility and marriage across time, and heterogeneity in the population, using data on all Norwegian cohorts since 1951 who remained present in the country at age 40. We make use of data on their outcomes from 1967 onwards. In particular, we are interested in how the probability of being childless, total fertility, and the probability of being partnered varies with the man's relative within-cohort earnings rank, and how these patterns have changed over recent decades. We then explore potential mechanisms by studying how relative earnings have changed over time, and how other outcomes such as health and incarceration vary with relative earnings.

Our measure of lifetime earnings rank draws on annual earnings from work covering the period 1967 to 2018. To bypass the need for deflation, for each individual we first computed the within gender and birth cohort earnings percentile at each adult age. Next, we took the average of these percentiles over the age span 30 to 60, and recomputed the individual's lifetime earnings rank from the distribution of average percentiles. We use this measure of within-gender lifetime earnings rank to characterize our study population of men, as well as their fathers.[4] Specifically, we show how outcomes vary with own lifetime earnings rank and father's lifetime earnings rank. We include father's lifetime earnings rank as a comparison point, because it removes some of the endogeneity associated with own income, and may better capture an individual's ex ante lifetime opportunities.

3.1 Two Facts on Male Childlessness and Total Fertility

Fact 1: Male childlessness is highest among men with lowest 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 cohorts. The pattern is striking: while only around 10% of men in the top 5% of the own earnings distribution are childless, this number jumps to around 60% in the bottom 5%. In the most recent cohort, 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 exponentially below the 30th percentile of the earnings distribution. In the empirical analysis of bankruptcies, our sample consists of lower earning men compared to the population mean, thus aligning with the steeper fertility gradient in earnings.

[4]The algorithm allows us to use the full 31 years of age-specific percentiles for those born between 1937 and 1958. For the youngest cohort of our study population (born 1979), the rank measure is based on ten observations covering the age span 30 to 39. Conversely, 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).

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 men's own earnings rank is more predictive of childlessness than father's rank.

Panels C and D depict the relationship 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.

Figure 1: Fertility across the earnings distribution.

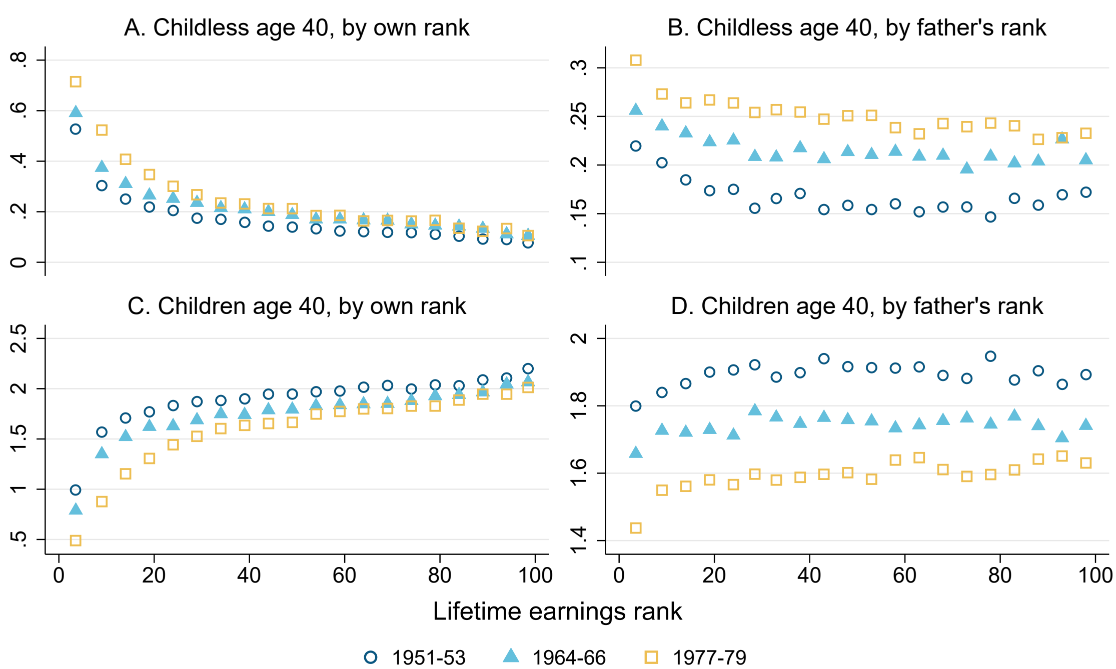

Notes: Each scatter point represents five percent of Norwegian men born between 1951-1953, 1964-1966, and 1977-1979, respectively. Observation count is 245 113.

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 representative cohorts, 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. This shows a striking fact: the difference between childlessness rates at the bottom and top of the earnings distribution has widened substantially over time. While the 1951 cohort had a range of 35 percentage points, this widened to 51 percentage points for the 1979 cohort. We still see that childlessness rates are highest for those in the lowest ranks, and that childlessness rates overall have increased over time. These relationships are less pronounced for father's earnings rank, but 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, 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 0.88 children for men born in 1951 to 1.34 children for men born in 1979.

We conduct an additional exercise to check that these trends are not driven by increasing delay in having a child. In the Appendix, we show comparable figures with fertility at age 50, rather than age 40 (Figure A.1). Note that this restricts the number of cohort years we can represent in Figure A.1, with 1969 being the youngest cohort present in both the fertility at age 40 and fertility at age 50 figures. Comparing like-for-like cohorts, we see that the patterns are very similar in both figures, with similar rates and gaps in childlessness and number of children across the earnings distribution, over time. The increase in rates of childlessness among the lowest earners is particularly striking and robust, and there is little evidence to suggest that the patterns can be explained by increased fertility delay.

3.2 Marital Status and Number of Partners

As a complement to the stylized facts on fertility, it is natural to consider whether these patterns are reflected in marital status and number of partners. In particular, it may be that these relationships are driven by the marriage market, with the lowest earning men being unable to find partners and therefore to have children. On the other hand, the effect may be driven by what would have been out-of-wedlock births, and therefore the relationship between earnings rank and marital status may be more muted. Figure 3 sheds light on this question. In Panel A, we plot the average proportion of men who are single (neither married nor cohabiting) by their position in the lifetime earnings distribution, for the three birth cohorts in the beginning, middle and end of our observation period. There are some 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.

This naturally leads to the question whether male fertility has been concentrated among those with better labor market prospects via partnership. In particular, are the best men being "recycled" and having children with multiple women? Panel C shows that this is indeed the case, with the highest number of partners by age 40 seen for men at the top of the earnings

Figure 2: Inequality in fertility over time.

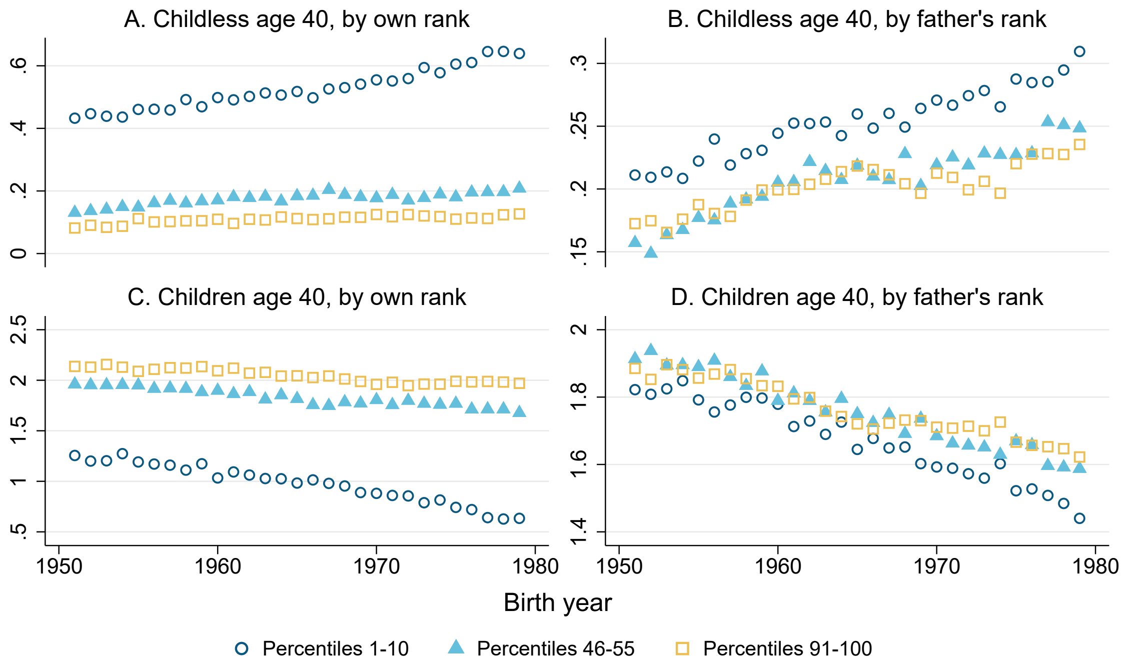

Notes: Scatter points represent ten percent of each cohort of Norwegian men born between 1951 and 1979.

distribution. This gap in the average number of partners by age 40 between the lowest and highest earning men has also widened over time, indicating that the marriage market is an important component of the fertility-earnings relationship. To this end, we also explore the impact of job loss via bankruptcies on partnering in Section 5.

Figure 3: Marital status and number of partners.

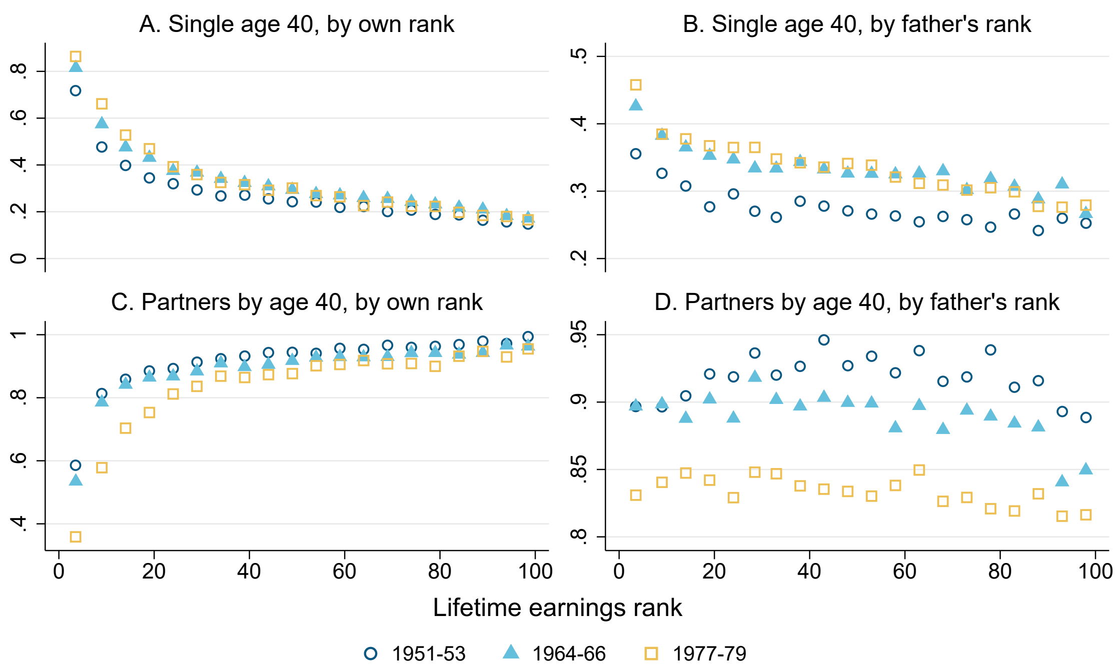

Notes: Each scatter point represents five percent of Norwegian men born between 1951-1953, 1964-1966, and 1977-1979, respectively. Panels C and D count the the number of unique partners with whom the male has fathered a child, including zero for those childless at 40. Observation count is 245 113.

3.3 Potential Mechanisms: Relative Earnings, Health, Incarceration, and Data Quality

There are two crucial stylized facts that emerge from an analysis of the relationship between fertility and earnings rank: childlessness rates are highest for the least well-off men, and this inequality has increased over time, with the gap in childlessness rates between men at the top and bottom of the earnings distribution widening over time. We have also documented that the relationship between single status and earnings rank is similar, suggesting that a key mechanism for these relationships may be economic returns of men on the marriage market. To explore this further, it is instructive to analyze how relative earnings have changed over time, as well as other potential outcomes that can correlate with both earnings and fertility: health, incarceration and missing data.

Relative Earnings Thus far we have considered the relationship between fertility and relative earnings rank, but this does not shed light on how the earnings of those at the bottom of the distribution have changed over time. Figure 4 depicts real absolute earnings at age 40 in 100.000 NOK by cohort, for the three points of the earnings distribution. 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 12 times the earnings of men in the bottom 10%, as compared to a multiple of 6 for the earliest cohort in the figure. A similar though less pronounced pattern is seen by father's earnings rank. Insofar as labor market earnings are a determinant of returns on the marriage market, this suggests that the marriage market value of men at the lower end of the earnings distribution has declined over time, in relative terms, and is consistent with the patterns of childlessness and partnering we have seen above.

Figure 4: Absolute earnings over time.

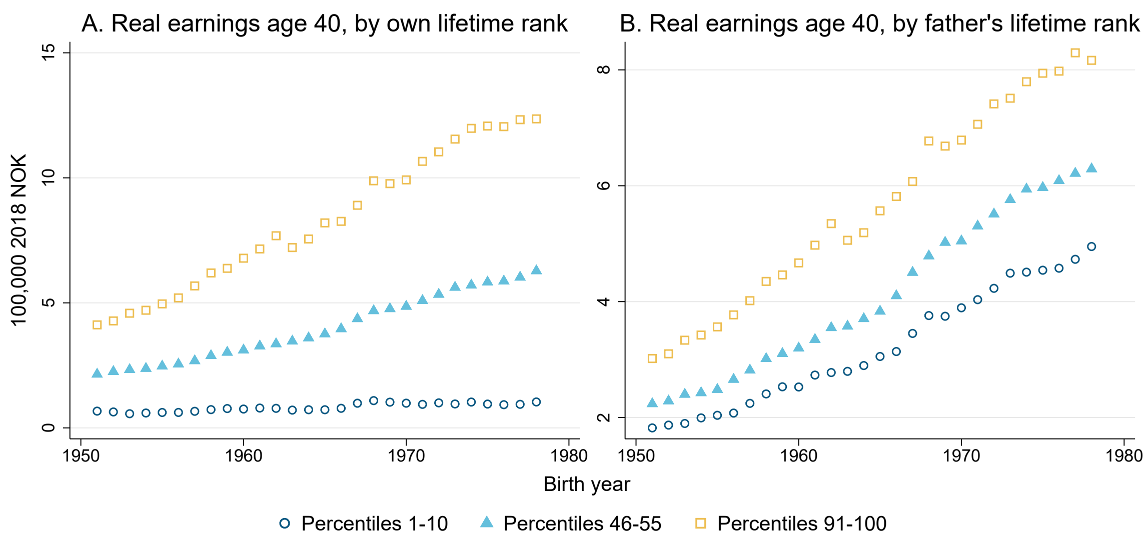

Notes: Scatter points represent ten percent of each cohort of Norwegian men born between 1951 and 1978. Earnings are observed at age 40, are inflated to 2019 NOK, and are depicted in units of 100 000. Observation count is 390 784.

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. 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 5 depicts the relationship between relative earnings rank and the average proportion of individuals registered as having a long-term disability at age 30.[5] Although there is a negative correlation between relative earnings rank

[5]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.

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 5: Disability and earnings.

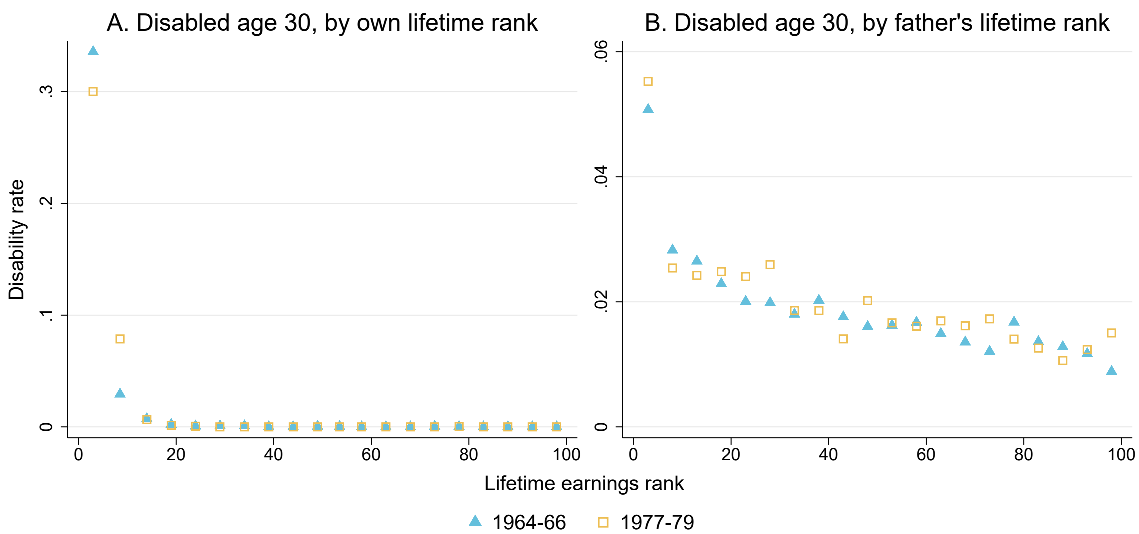

Notes: Each scatter point represents five percent of Norwegian men born between 1964-1966 and 1977-1979, respectively. Disability status is measured by receipt of a permanent disability pension at age 30. Observation count is 162 412. The average disability rate is 0.020.

Figure 6 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 7 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.[6] 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.

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

Figure 6: Earnings and other markers of health.

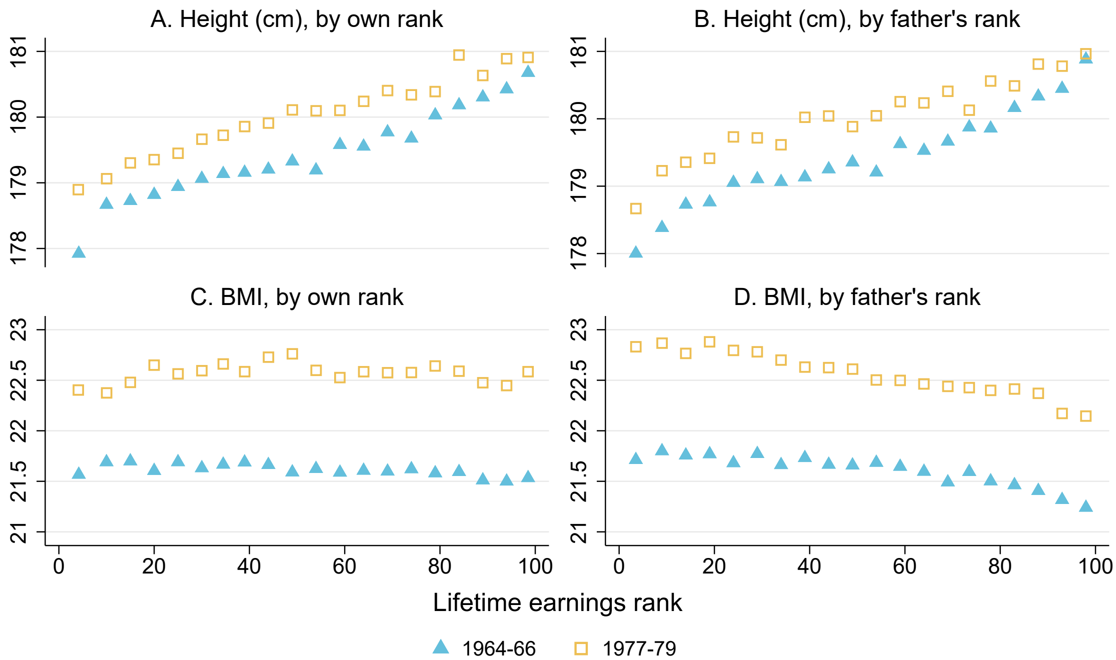

Notes: Each scatter point represents five percent of Norwegian men born between 1964-1966 and 1977-1979, respectively (data for the 1951-1953 cohorts are not available). Height and weight are measured at conscription for military service, typically at age 17 or 18. Observation count is 158 531. The average height is 179.7 cm and average BMI is 22.0.

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 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 8 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.[7] 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.

[7]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 7: Incarceration and earnings.

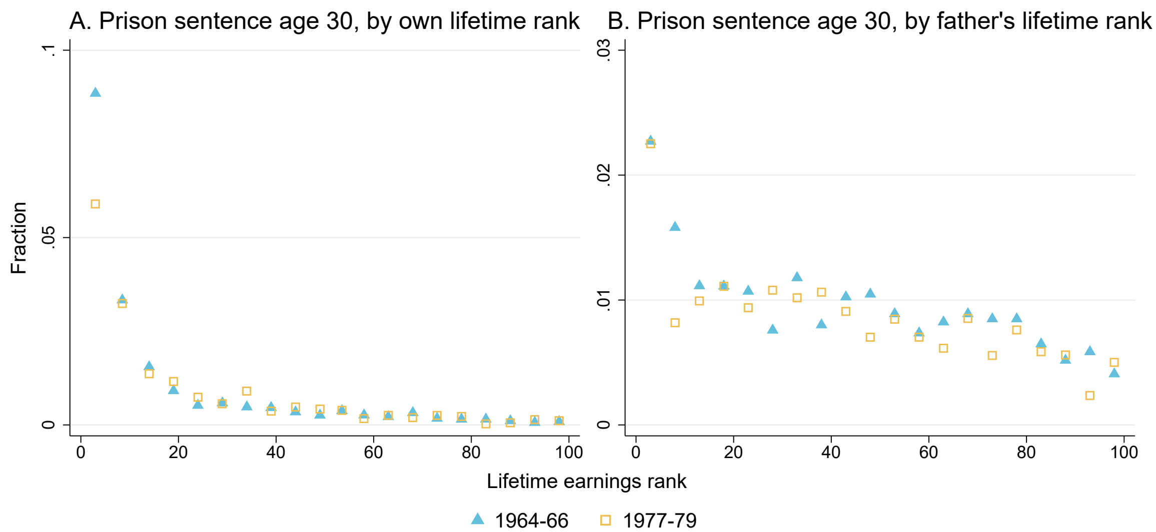

Notes: Each scatter point represents five percent of Norwegian men born between 1964-1966 and 1977-1979, respectively. Scatter points give the fraction of men charged with a crime and sentenced to unconditional imprisonment the year they turned 30. Observation count is 162 412. The average imprisonment rate is 0.0096.

3.4 Evolution of inequality in fertility for women

Although not the main focus of the paper, it is also instructive to analyze these same patterns for women and we show these results in Appendix A.2. We document several interesting patterns. The relationship between women's childlessness rates and relative earnings rank is U-shaped: rates are highest at the extreme ends of the earnings distribution. This is consistent with the findings of Baudin, de la Croix, and Gobbi (2015) for the U.S., 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. Still, rates of childlessness do not vary across the earnings distribution for women nearly as much as they do for men. Turning to the relationship with father's rank, the pattern is almost flat: father's earnings rank seems to have little bearing on a woman's childlessness status. With respect to total fertility, the relationship with own earnings is U-shaped for women. The highest fertility rates are observed for women in the 15th percentile of own earnings.

Interestingly, regarding the evolution of inequality over time, the relationship between own earnings rank and fertility exhibits a crossing pattern, with the childlessness rates of women in the bottom of the earnings distribution increasing steadily with each cohort, while the childlessness rates of women in the top earnings rank decreasing over time. In this sense,

Figure 8: Missing birth records and mothers' earnings.

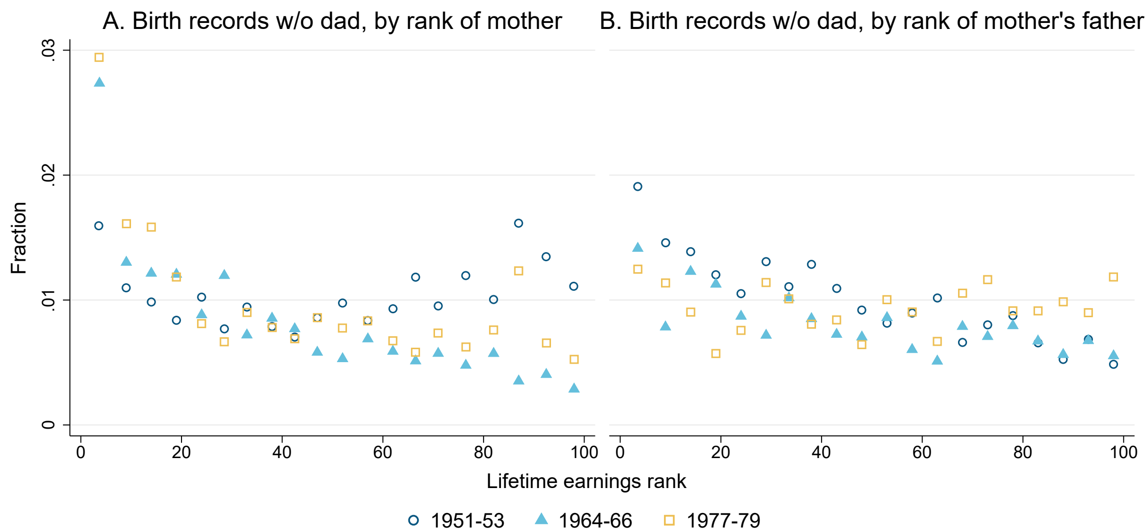

Notes: Each scatter point represents five percent of Norwegian women born between 1951-1953, 1964-1966 and 1977-1979, respectively. Scatter points give the fraction of birth records with missing information on the child's father. Observation count is 472 794 children born to 206 935 women by age 40. Average rate is 0.0066 per birth record.

the family penalty to "career women" has declined over time. Childlessness rates for women in the middle of the earnings distribution have remained stable and low over time. The relationship with father's earnings rank is weak, but suggestive of a similar pattern for men: over time, childlessness rates for women whose fathers had a higher earnings rank have fallen, and for women whose fathers had a lower earnings rank have risen. With respect to total fertility for women we see a crossing between the bottom 10% and top 10% earnings groups at around the 1975 birth cohort.

Earnings inequality has also widened for women, but the gap between the top and bottom percentiles of women's earnings is much smaller than for men.

4 Empirical Strategy to Identify Effects of Earnings on Male Fertility

We have documented a striking inequality in male fertility across the earnings distribution. These correlations may, however, 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. We use bankruptcies to identify the impact of earnings losses on labor market and, most importantly, family outcomes. Using our descriptive analysis as a jumping-off point, we estimate the impact of bankruptcies on the probability of having a child in a given year, cumulative fertility, having a partner, and a comprehensive set of labor market outcomes to verify our first stage. We also check the impact on disability status as an alternative mechanism, and present an extensive set of robustness checks at the end of the Section that check for issues such as sampling, selection, heterogeneous treatment effects and underlying trends.

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). They are relatively common, with 1% of the Norwegian working population experiencing a bankruptcy in any two years. We verify that they 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 firm closures in general 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.

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. Our approach makes use of within-individual time variation in exposure to the shock. Following Lindo (2010), we include individual fixed effects to account for any time-invariant characteristics that may affect both exposure to bankruptcy and family outcomes. This means that our estimates will be unbiased even in the presence of time-invariant unobservable characteristics that correlate with both exposure to bankruptcy and the set of outcomes. To assuage concerns over bias arising from time-varying unobservable characteristics, we include a control group of matched same-sex siblings working in a firm that does not go bankrupt. They 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.[8] Still, we dig further into such

[8]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.

possible differential trends in Section 5.3. We 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, and conduct a stacked regression design to allow for heterogeneous treatment effects over time. We find that our conclusions are robust to these checks.

Taking together the bankruptcy event study and the control group of same sex siblings yields the following estimating equation:

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

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. We estimate impacts from seven years before to 15 years after the bankruptcy; our long pre-bankruptcy window allows us to check potential differential trends over a longer time period.[9] 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.

Bankruptcies may be anticipated, and individuals with better outside options, and differential family outcomes, may leave before losing their job and be missing from our sample. Alternatively, firms in distress may have slower wage growth than non-distressed firms. These selection and compositional concerns are discussed in Dustmann and Meghir (2005),

[9]Not all data are available for all outcome years, but 58 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 5.3).

who consider sampling individuals either one year or two years prior to the firm closure. We 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 5.3, by changing the timing of when we sample individuals.

We set the omitted year to be -5. Therefore, 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. As we discuss in the results section, and consistent with Dustmann and Meghir (2005), there is evidence of some selection in the years prior to treatment, but we see no differential trends at five years or earlier. Therefore, we argue that year -5 provides the cleanest measure of pre-treatment outcome levels. 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).

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).

5 Effects of Firm Bankruptcies

5.1 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). Recall that the earnings-fertility gradient documented in Section 3 is steepest among the lowest earners and therefore it is useful to focus on the family and labor market outcomes of lower earners in the event study sample. 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 is less of a problem for us since we include 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 9 and 11 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 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 (1)	Control (non-bankrupt) firms (2)
Observations	262 865	311 608
Individuals	13 087	16 121
Firms	6 873	8 581
Mean firm size	44.2 [115.0]	842.5 [2381.6]
Mean firm age	9.6 [8.0]	17.2 [13.2]
Mean firm log wage	5.164 [0.339]	5.348 [0.305]
Manufacturing	0.226 [0.418]	0.197 [0.398]
Construction	0.229 [0.420]	0.164 [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 [0.150]
Info/communications	0.050 [0.218]	0.046 [0.211]
Prof/tech services	0.042 [0.200]	0.039 [0.194]
Admin/support services	0.048 [0.215]	0.047 [0.213]
Public admin	0.000 [.]	0.052 [0.222]
Health services	0.009 [0.097]	0.051 [0.221]
Other	0.061 [0.240]	0.122 [0.327]

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 2019 NOK. Numbers in brackets are standard deviations.

5.2 Results

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.

Labor market outcomes As an initial analysis into how labor market outcomes evolve before and after bankruptcy, Figure 9 compares the means over time for the men exposed to bankruptcies and their matched brothers. These are sample means that do not account for any control variables. 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 10 depicts 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.[10] 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. While there is little noticeable prior diverging trend in being registered for a November job, however, we note that being registered as unemployed at any time in the year (Panel B) already increases around three years before the bankruptcy. This is consistent with Dustmann and Meghir (2005) who show that the impacts of a firm closure may already be evident two years before the event.

Panel C shows a substantial decline in the number of hours worked for men exposed to bankruptcy, as well as a decline in hourly pay (Panel E). These are important findings because the sample for hours worked and hourly pay is conditional: it only consists of men actually working, whether at the bankruptcy firm or at a different firm (for example, if they left and found a better job). They are likely to be positively selected, with the lower productivity men laid off first and unable to find a replacement job. Thus, a declining prior trend in working hours and hourly wage is highly likely to be driven by firm-specific trends, rather than individual ones. We also show estimated impacts on an unconditional sample of hourly wages and total working hours (Panels D and F), where we predict these outcomes for those men with missing data (the unemployed). The impacts are similar but more negative, consistent with our argument that the sample in Panels C and E is one of positively selected

[10] 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.

men.

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 income.

An advantage of our data and setup is that we are able to analyze a longer pre-treatment window. It becomes clear that the 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 lower also at year -4). We find this largely reassuring as although some differences are to be expected when comparing men working in distressed and non-distressed firms (Dustmann and Meghir 2005), these differences do not appear to be life-long. Instead, they are more likely to be explained by characteristics of the firm, such as declining or stagnant wages in struggling firms over the space of three to four years preceding bankruptcy (Jacobson, LaLonde, and Sullivan 1993). Still, we take extra care with exploring alternative control groups and other checks on pre-treatment trends in Section 5.3.

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).

Fertility outcomes In Figures 11 and 12, we depict the control-treatment means and estimated coefficients from Equation (1) on fertility and family 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. The estimates in these Figures show a similar pattern. There is a divergence in average fertility and family outcomes after the bankruptcy event, which is confirmed in the event study estimated coefficients: the probability of being unpartnered increases significantly following exposure to a firm bankruptcy, reaching a peak of 3.4 percentage points in year 4, with little relationship seen before the event (Panel A). Men exposed to a 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 of experiencing a first birth - transitioning out of childlessness - is also lower, and makes up more than half of the effect of Panel B. The impact of higher parity births is more muted and somewhat delayed compared to first births (Panel D). Finally, the effect on cumulative fertility is negative and grows over time (Panels E and F). Up to ten years after the bankruptcy event, men who experienced this negative labor market shock are significantly more likely to be childless, and the negative impact on total fertility is remarkably persistent and does not recover during our sample window of 15 years.

This indicates that the effect on total fertility stems from “missed births” in the initial few years after job loss, that are not compensated for in later life. This is likely to stem from both the reduced rate of partnering in the initial years after the bankruptcy shock (Panel A), as well as reduced fecundity with age for those men with partners who choose to postpone having a child due to the labor market shock.

Importantly, prior to bankruptcy, there is no discernible difference in fertility trends between the bankrupt and non-bankrupt groups of men. We note some delay between impacts on labor market outcomes and fertility outcomes, which is to be expected given that fertility intentions take some time to materialize.

Difference-in-difference estimates To provide further evidence for the impacts of bankruptcy on labor market and fertility outcomes, we estimate a difference-in-difference regression comparing average outcomes before and after bankruptcy. Taking our event study estimates as guidance, we 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). The idea is to provide a cleaner comparison of the pre-treatment and post-treatment periods, removing the years where we see evidence of pre-trends in Figure 10. We find statistically significant 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. In column (2), we estimate an extended donut also omitting years -2 and -3 (as we note above declining trends in income and employment in those years). This 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.

Scaling 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 find that the probability of single status increases by 6.7% and childlessness by 3.9%, while the probability of experiencing the birth of a child and total fertility decline by 10.4% and 2.2% respectively.[11]

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. Taken together with our descriptive results on the inequality in family life across the earnings distribution, this suggests an important connection between labor market prospects and men's access to family life, with a particular vulnerability among the lower male earners in the population. In Section 6, we provide a back-of-the-envelope calculation

[11]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.

using our causal estimates to quantify the descriptive evidence, and take a wider lens to look at how the relationship between job loss and fertility has changed over the last three decades.

Figure 9: Sibling mean comparisons before and after firm bankruptcies, labor market outcomes.)

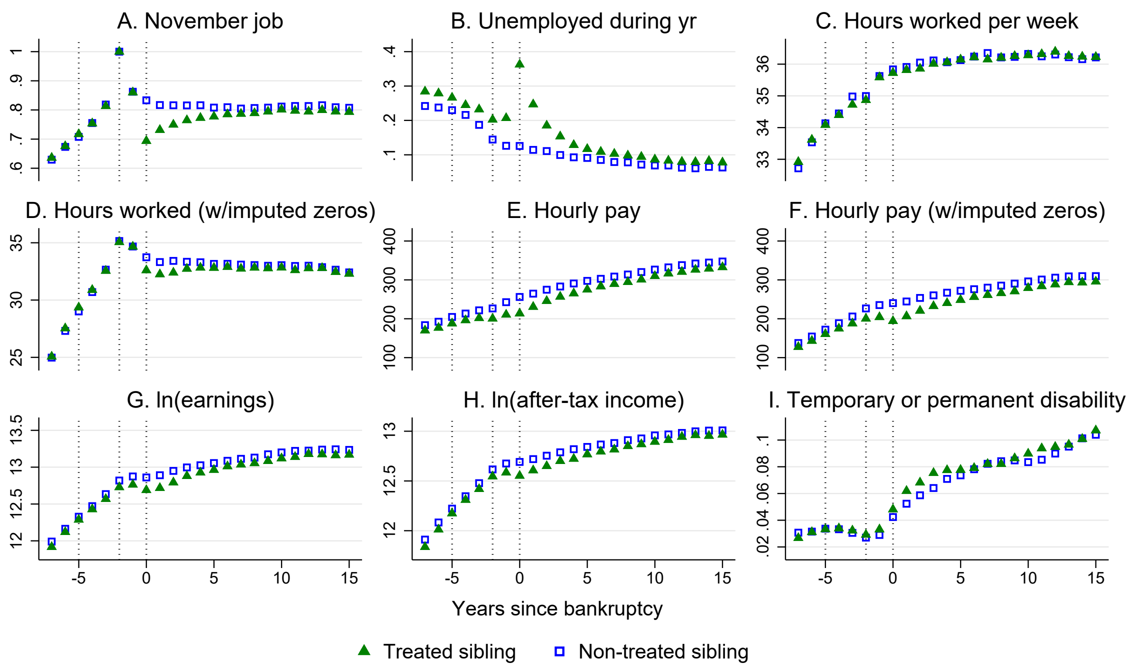

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Sample of treated siblings consists of Norwegian-born men who in year -2 were employed in a firm filing for bankruptcy two years later and age 25-35 in the year of the event, while 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 262 865 in the treatment group and 311 608 in the control group, but may be lower for some outcomes with missing data in certain years. Wages and earnings are inflated to 2019 NOK.

Figure 10: Effects of firm bankruptcies on labor market outcomes.

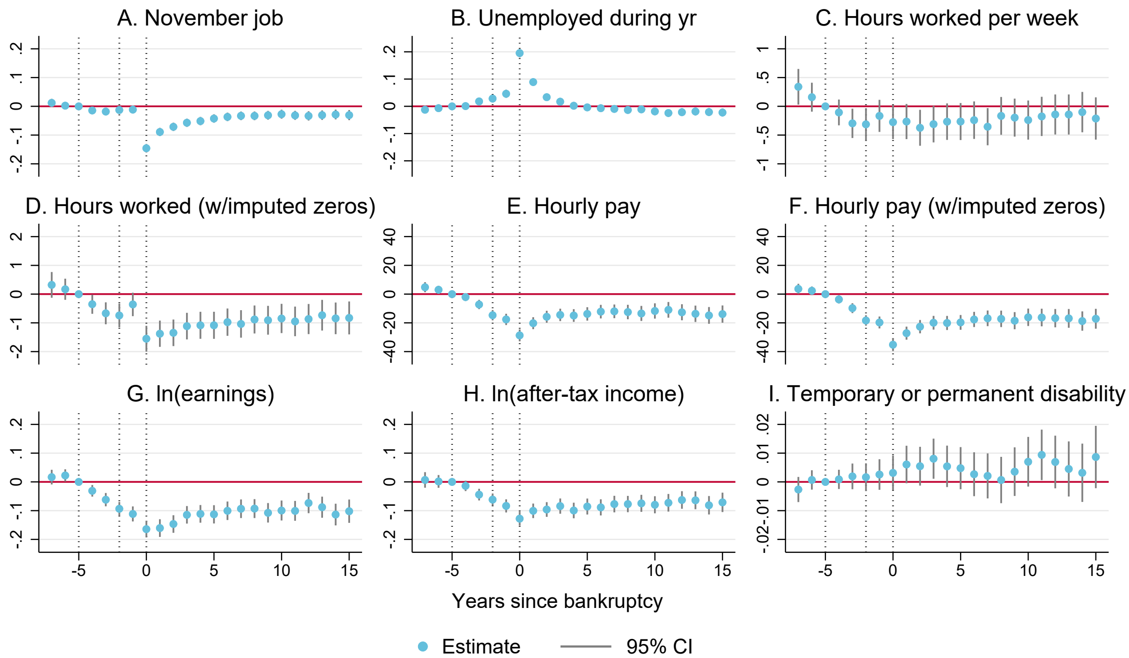

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Sample of treated siblings consists of Norwegian-born men who in year -2 were employed in a firm filing for bankruptcy two years later and age 25-35 in the year of the event, while 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 262 865 in the treatment group and 311 608 in the control group.

Figure 11: Sibling mean comparisons before and after firm bankruptcies, fertility outcomes.

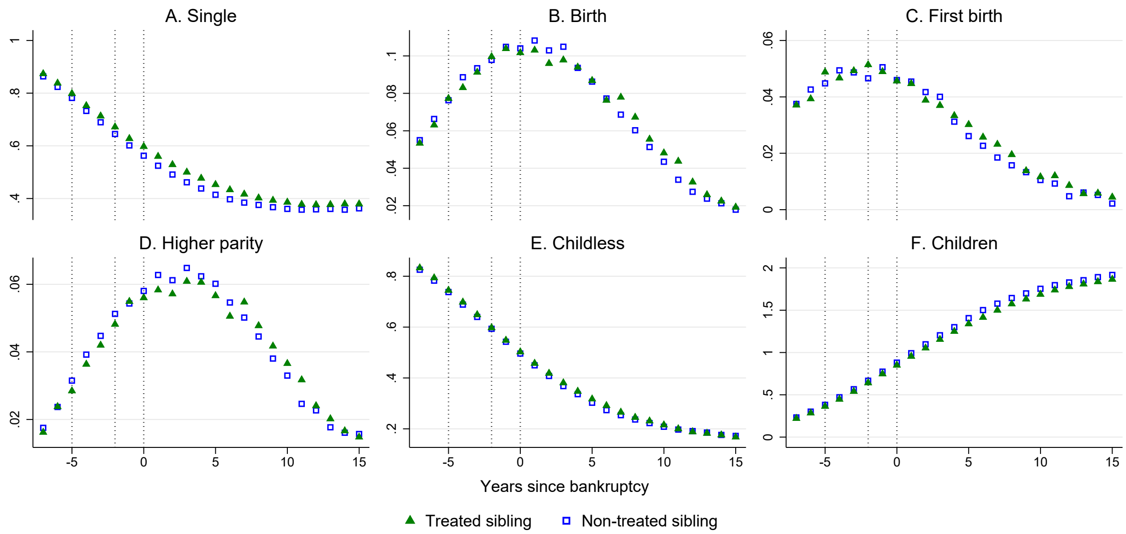

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Sample of treated siblings consists of Norwegian-born men who in year -2 were employed in a firm filing for bankruptcy two years later and age 25-35 in the year of the event, while 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 12: Effects of firm bankruptcies on fertility outcomes.

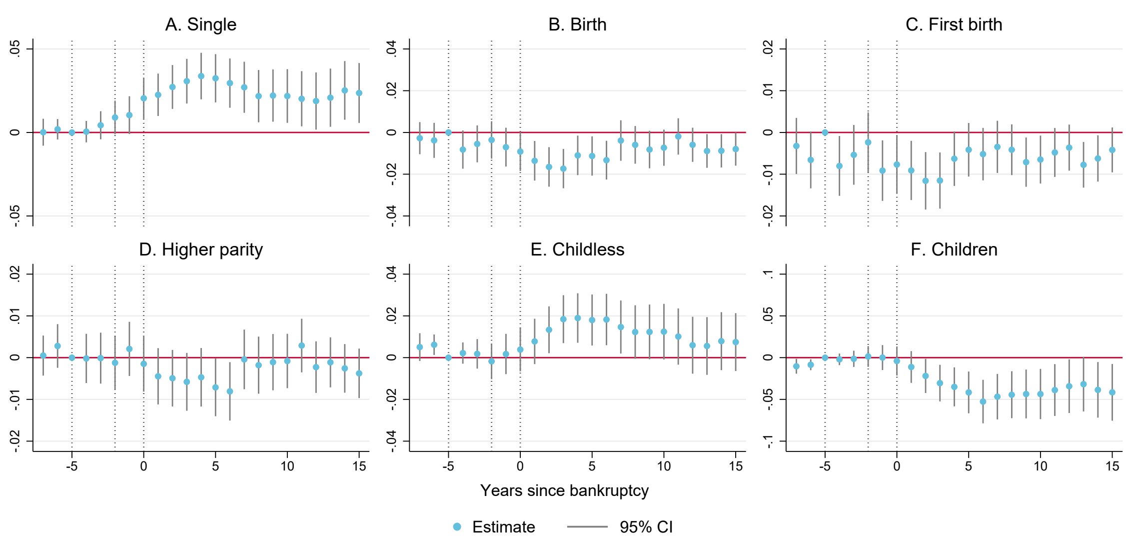

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Scatter points show the estimates of β_t from the estimating equation. See text and notes to 11 for a description of the samples.

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	-0.045***	-0.047***
	(0.005)	(0.005)
Unemployed during yr	-0.001	0.009**
	(0.004)	(0.004)
Hours per week	-0.878***	-1.090***
	(0.167)	(0.199)
Hourly pay	-9.420***	-14.765***
	(1.687)	(2.048)
ln(earnings)	-0.092***	-0.116***
	(0.011)	(0.013)
ln(after-tax income)	-0.066***	-0.084***
	(0.009)	(0.011)
Fertility outcomes		
Single	0.025***	0.028***
	(0.006)	(0.006)
Birth	-0.006***	-0.007***
	(0.002)	(0.002)
First birth	-0.003**	-0.002*
	(0.001)	(0.001)
Higher parity birth	-0.004***	-0.005***
	(0.001)	(0.001)
Childless	0.012**	0.011***
	(0.005)	(0.005)
Children	-0.034***	-0.033***
	(0.011)	(0.012)
Observations[1]	516 400	458 149
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. All outcomes are without imputation. 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 479 532.

5.3 Robustness Checks

We conduct important checks on sampling, heterogeneous treatment effects, pre-existing trends and employment definitions in this section.

Time paths of outcomes in different samples 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). Another way of thinking about this choice is that it is a sample of individuals who have not yet left the firm. This may lead to some selection on outcome variables, which we explore 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 13.

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 13: Evolution of outcomes over time for different samples.

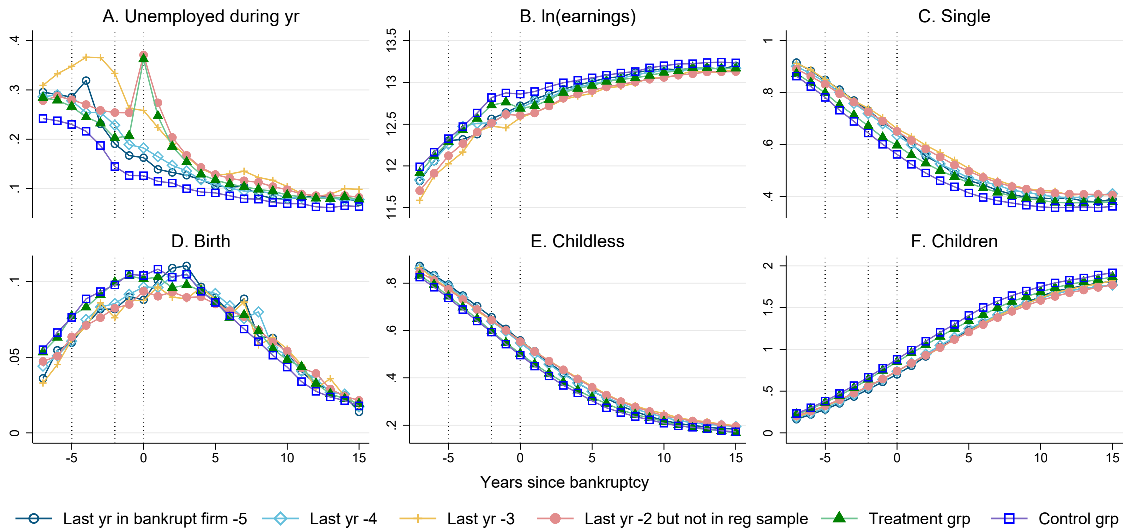

Notes: Samples consist of men age 25-35 at time 0, separated by the last year of employment at the firm filing for bankruptcy (at time 0). For completeness, the figure adds the time paths for the relevant outcomes for treatment and control groups depicted in Figures 9 and 11.

Choice of sampling year To complement our analysis of alternative samples, in this section we report estimates where we 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 14 show the results (Figure A.2 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 and first births is less marked here, with negative coefficients that are not statistically significant. This is to be expected given that we are introducing more measurement error by having a less precise treatment sample.

Figure 14: Effects of firm bankruptcies, sampling at t-3.

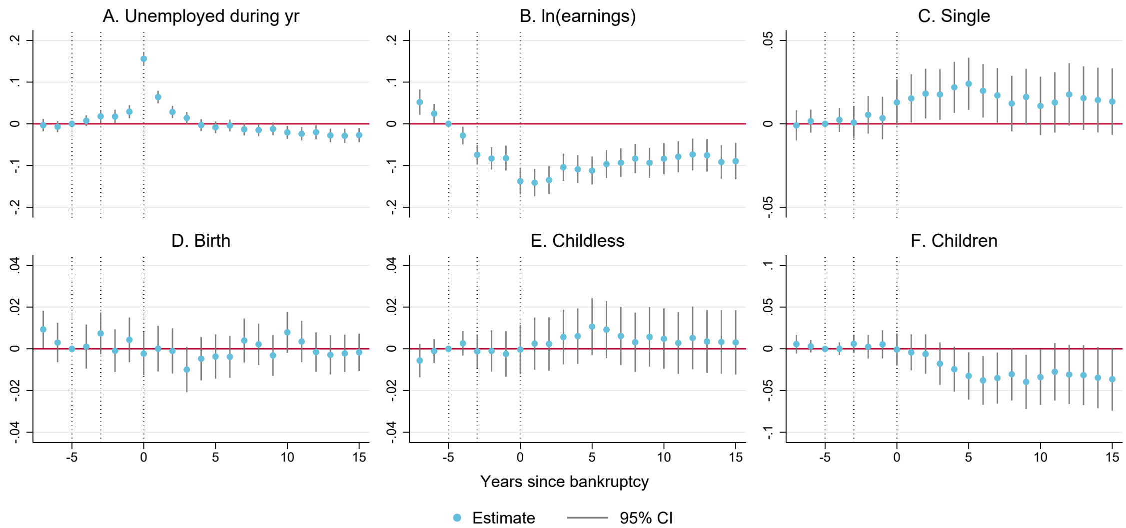

Notes: Vertical lines indicate year of observed November job (year -3), year of event (year 0), and reference year (-5). Scatter points show the estimates of β_t from the estimating equation. Sample of treated siblings consists of Norwegian-born men who in year -3 worked in a firm that filed for bankruptcy 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 209 233 in the treatment group and 246 900 in the control group.

Alternative definition of workplace closure We also 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) and Huttunen, Møen, and Salvanes (2011), but yields a more broad definition of workplace closure and although we minimise false shutdowns, we may not be able to rule them out entirely, which can introduce measurement error. Moreover, the closure of an establishment likely represents a less abrupt change compared to a firm bankruptcy. Figure A.3 in the Appendix shows mean outcomes over time for the two comparison groups. Figure 15 shows that, although our main estimated effects on unemployment and earnings persist here, they are about one half the magnitude of those in Figure 10. 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. This may not be surprising as our sample now includes all workplace closures; these may be more easily anticipated than those following a bankruptcy.

Figure 15: Effects of establishment shutdowns.

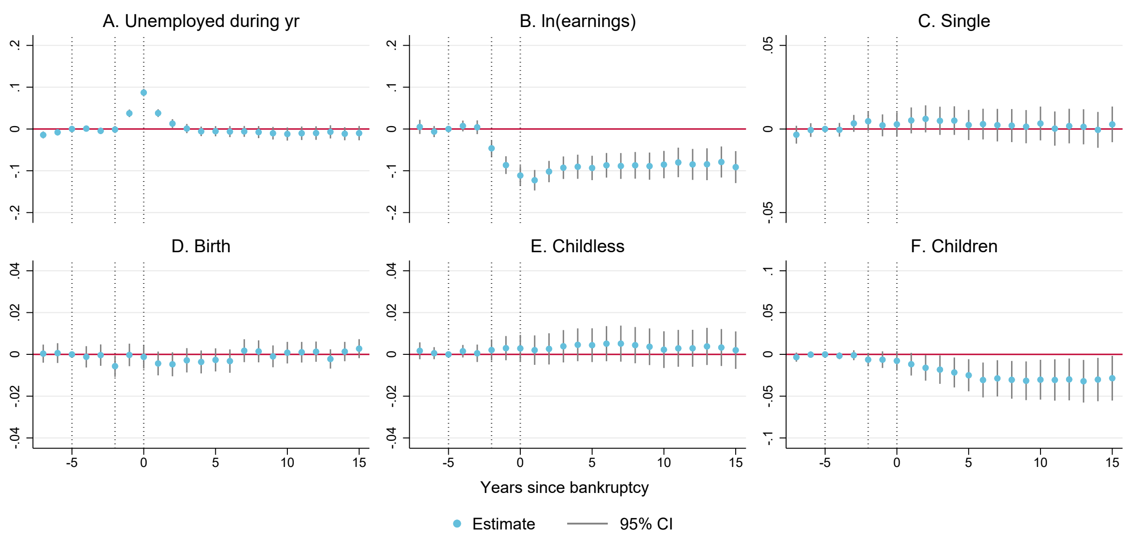

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Scatter points show the estimates of β_t from the estimating equation. Sample of treated siblings consists of Norwegian-born men who in year -2 worked at an establishment that shut 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 791 001 in the treatment group and 932 514 in the control group.

Allowing for year specific family FEs As an alternative check, 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. This is a valuable additional check to account for possible pre-trends. Reassuringly, Figure 16 shows that the estimates are essentially unchanged.

Figure 16: Effects of firm bankruptcies, accounting for family-by-year fixed effects.

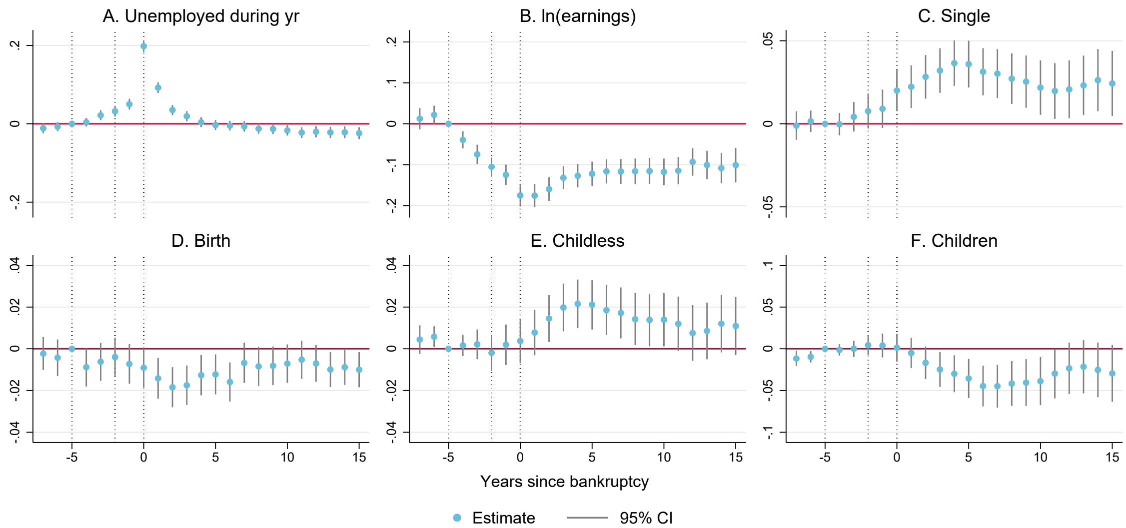

Notes: Regression model is augmented with family-by-year fixed effects. See also note to Figure 10.

Removing bankruptcies in other years 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 17 shows the estimates, which are not sensitive to this stricter sample restriction.

Figure 17: Removing any alternative bankruptcy events.

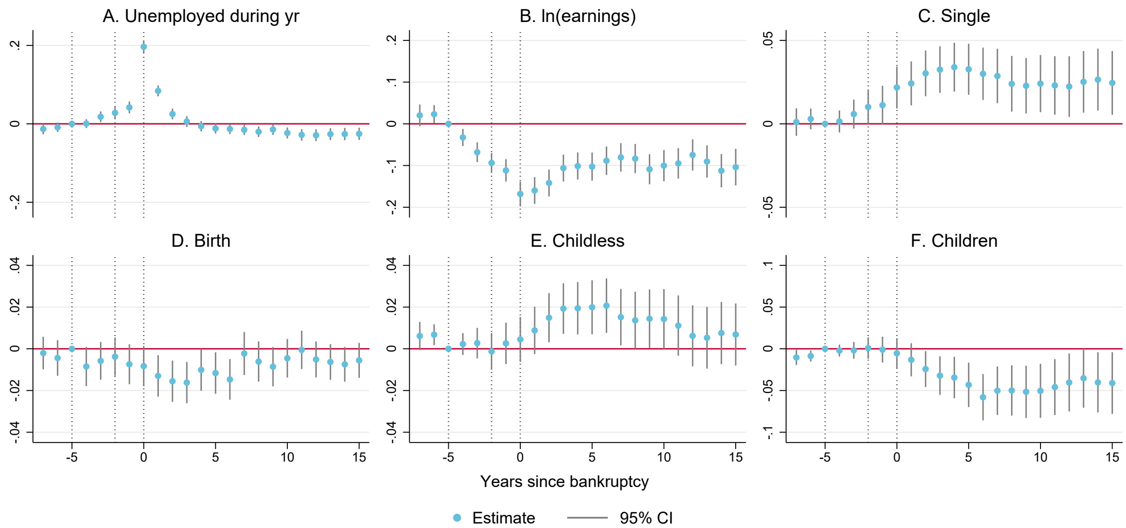

Notes: Regression samples exclude individuals who experience bankruptcy in other years than yr 0, so that the treatment sample is restricted to individuals who experience only one bankruptcy and the control sample to individuals who do not experience a bankruptcy during the observation window. Observation counts are 227 630 in the treatment group and 291 820 in the control group.

Balanced sample 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 1995-2004 (compared to the baseline 1995-2015) whose fertility outcomes can be tracked for the full 23 years. Civil status is first available in 1992, however, while data on earnings end in 2018. 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 18, with means reported in Figure A.4 in the Appendix).

Figure 18: Effects of firm bankruptcies, balanced sample.

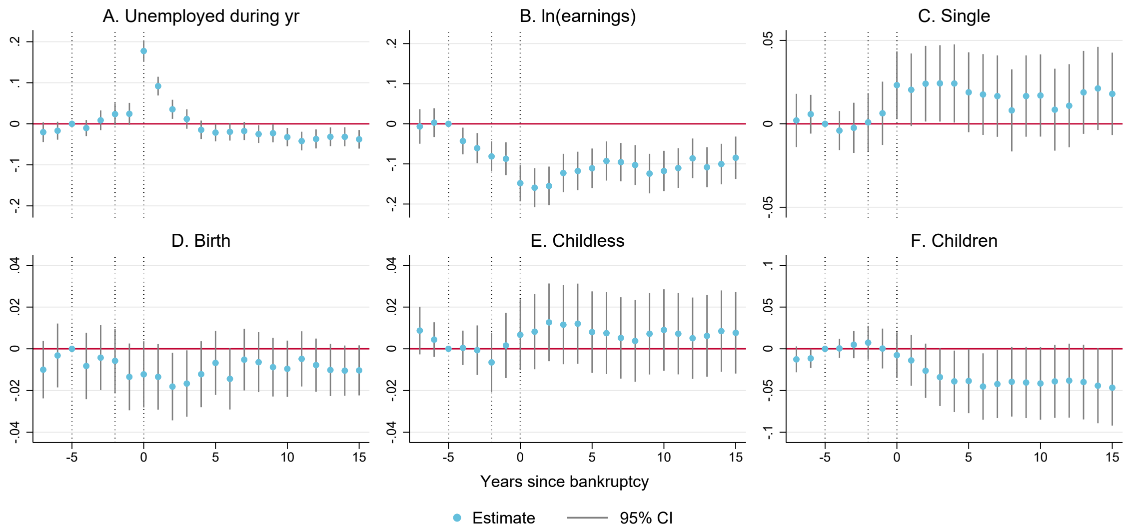

Notes: Regression model is estimated on restricted balanced sample where the 23-year time sequence falls within the data window 1988-2019, with bankruptcies in 1995-2004. Observation counts are 112 397 in the treatment group and 129 051 in the control group.

Stacked event-by-event analysis 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 9) 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 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 19, 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.

Figure 19: Comparing our main estimates to estimates from a stacked regression approach.

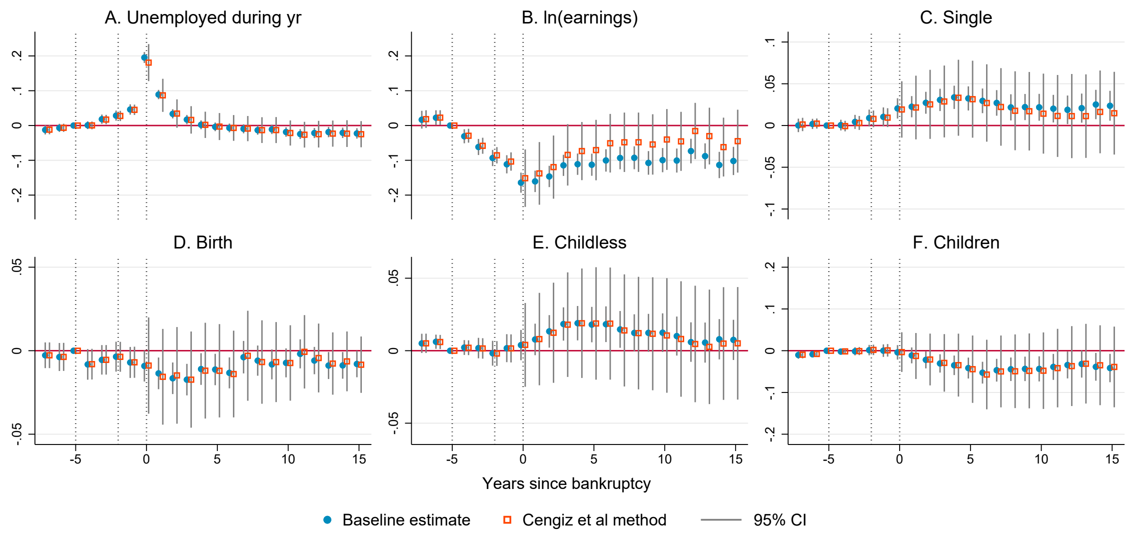

Notes: Baseline estimates replicate those in Figures 10 and 12. The Cengiz et al approach draws on a stacked event-by-event analysis, where each point estimate is based on 21 separate regressions omitting any observations where already treated individuals may influence estimation of calendar year effects.

6 Labor outcomes and male fertility: What do we learn?

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. In this section, we bring the two exercises together to draw wider conclusions about the relationship between labor outcomes and male fertility. First, we conduct a back-of-the-envelope calculation that yields a measure of the share of the descriptive relationship between labor market outcomes and fertility that is likely to be causal. Second, we show that this relationship between labor market outcomes and fertility has changed over time, in line with the widening inequality in fertility outcomes shown in Section 3. Together, these exercises add to the evidence that poor labor market outcomes have a negative impact on men's fertility outcomes, and that this impact has worsened over time.

6.1 Linking the descriptive and causal estimates

We conduct a back-of-the-envelope exercise to estimate the share of the descriptive relationship between fertility and earnings that is likely to be causal. First, we 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). Figure 20, Panels A and D, displays 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 23.1 percentage points, and with having 0.61 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 0.033, while raising the likelihood of childlessness by 1.1 percentage points. Scaling this up to 1 unit of log earnings yields magnitudes of 0.28 children and 9.5 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, 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 46% (0.28 / 0.61 * 100) of the descriptive relationship between earnings and total fertility and 41% (9.5 / 23.1 * 100) of the comparable relationship between earnings and childlessness can be explained by a causal relation. This brings an added layer of evidence to the dramatic patterns between male fertility and labor market prospects that we have documented.

Figure 20: Binned scatter plots, fertility outcomes age 40 and real earnings at ages 21-40.

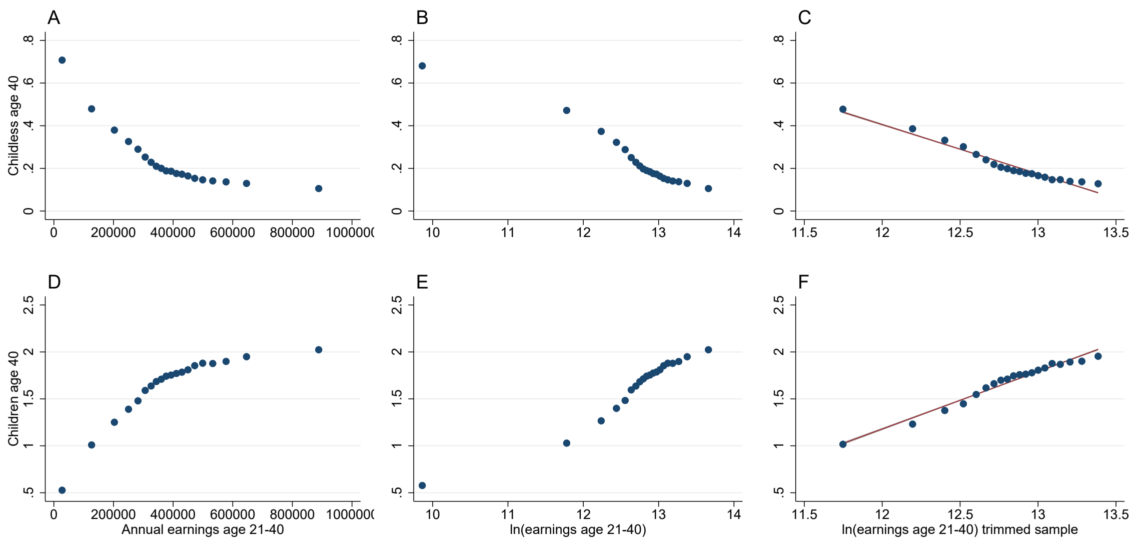

Notes: Sample consists of men born between 1971 and 1979. Earnings are inflated to 2019 NOK. Sample in Panels C and F is restricted to 5th to 95th percentile range of ln(earnings). Slopes (se) of regression line are -0.231 (0.002) in Panel C and 0.612 (0.006) in Panel F. Shaded area around regression line depicts 95 percent CI of prediction. Observation counts are 234 352 in Panels A and D, 233 438 in Panels B and E, and 210 086 in Panels C and F.

6.2 The Changing Relationship between Unemployment and Fertility over Time

Next, we provide additional evidence on the changing relationship between unemployment and male fertility over time. In Section 3, we documented the widening inequality in men's access to family life between low and high earners. Our findings using firm bankruptcies show that earnings losses and unemployment are associated with lower fertility and higher childlessness, but do not speak to the change in this relationship over time.

In order to investigate whether the relationship between job loss and fertility has changed over time, and whether this can plausibly explain the important facts uncovered in Section 3, we conduct the following exercise. We are interested in whether the *penalty* to job loss, in terms of fertility, has increased over time. Specifically, our bankruptcy analysis showed that job loss is associated with lower male fertility. Increasing inequality in male fertility can result from this if that impact has become more negative over time. This is what we explore in this Section.

Using cross-sectional population data for the period 1990-2019 for individuals aged 25-35 as in the event study sample, we regress the probability of experiencing the birth of a child 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:

$$\begin{aligned} Birth_{i,k,t} = \alpha + \sum_{t=1990}^{t=2019} \beta_t Year_{i,t} * Unemp_{i,k,t-1} + \gamma Exper_i + \tau Exper_i^2 \\ + \lambda Educ_i + \kappa_{i,k} + \theta_t Year_{i,t} + \eta_{i,k,t}, \end{aligned} \tag{2}$$

where $Birth_{i,k,t}$ indicates individual i having a child in year t living in municipality k, $Unemp_{i,k,t-1}$ 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 having a child this year and last year's unemployment status by calendar year.

In Figure 21, the top panels 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: being unemployed is associated with a higher probability of not experiencing the birth of a child in recent years, as compared to earlier 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. Panel C shows that the relationship for higher parity births is also negative, but with a less clear downward trend over time.

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 21 vary with the addition of controls related to this type of selection, such as education, BMI and IQ. Figure A.5 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.6, 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 21.

These striking findings show that job loss carries a higher penalty in terms of lower fertility 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. This provides additional evidence on the changing nature of men's family outcomes over time, and how they are affected by their labor market prospects. 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.

Figure 21: Unemployment and births over time

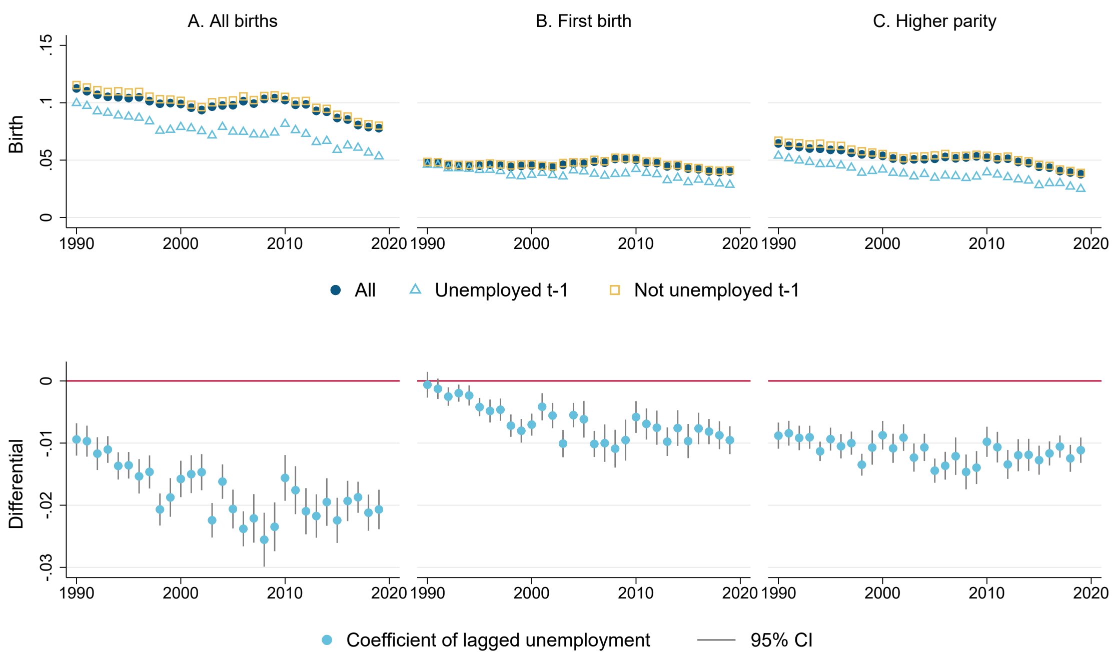

Notes: Scatter points in the top panels show fertility rates of men by unemployment status during the prior calendar year, while the bottom panels show the estimated coefficient of individual unemployment status from a regression of birth on registered unemployment the prior year. Regression controls 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 1990-2019. Observation count is 8 881 215. Mean birth rate is 0.098 and mean registered unemployment is 0.143.

7 Conclusion

Using detailed administrative data from Norway, we document a remarkable increase in the inequality of male childlessness across the income distribution. We further show that the poorest men are more likely to be single and that the income gradient in partnership formation has become steeper. To investigate whether labor market shocks may causally explain these descriptive facts, we use bankruptcies to identify the effect of job loss on fertility. We note significant and persistent negative impacts of bankruptcies on employment, earnings, births, total fertility and partnering rates. These do not recover for up to 15 years following the event. A simple calculation indicates that between 41%-46% of the descriptive earnings-fertility gradient may be driven by a causal relationship. We further show that the relationship between unemployment and fertility has become more negative over time, indicating stronger penalties in recent years for job loss in fertility.

Previous studies frequently do not have data on male fertility and those that do often investigate effects of job loss on fertility within existing couples, finding limited effects of male job loss. We argue that our estimates capture a wider set of effects of job loss on fertility as we include all men, even those single at the time of the shock. We find that bankruptcies affect partnering. As such, the total ramifications of job losses are not captured when conditioning on having a partner. Further, our data encompasses an entire population and we combine a rich descriptive analysis with a robust empirical strategy to show striking new findings on inequality in family life among men.

More generally, we provide new evidence for the existence of "left behind" men, who face wider consequences of stagnating earnings that reach beyond their labor market prospects. In addition, inequality may increase even further as there is a clear 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 also have wider societal ramifications that are not captured by focusing on earnings changes only: there is a shift in the distribution of new births in the population, which is likely to be accompanied by changes in child investments and quality.

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*.

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*.

Bhalotra, S., A. Venkataramani, and S. Walther (2021). Fertility and labor market responses to reductions in mortality. *Mimeo*.

Binder, A. J. and J. Bound (2019, May). The declining labor market prospects of less-educated men. *Journal of Economic Perspectives 33*(2), 163–90.

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.

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.

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.

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.

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.

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. (2004). The long road to the fast track: Career and family. *Annals of the American Academy of Political and Social Science 596*, 20–35.

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.

Hornstein, A., P. Krusell, and G. L. Violante (2005). The effects of technical change on labor market inequalities. *CEPS Working Paper* (113).

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.

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.

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.

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).

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.

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

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.

APPENDIX: FOR ONLINE PUBLICATION

A.1 Additional figures

Figure A.1: Inequality in fertility over time, measured at age 50.

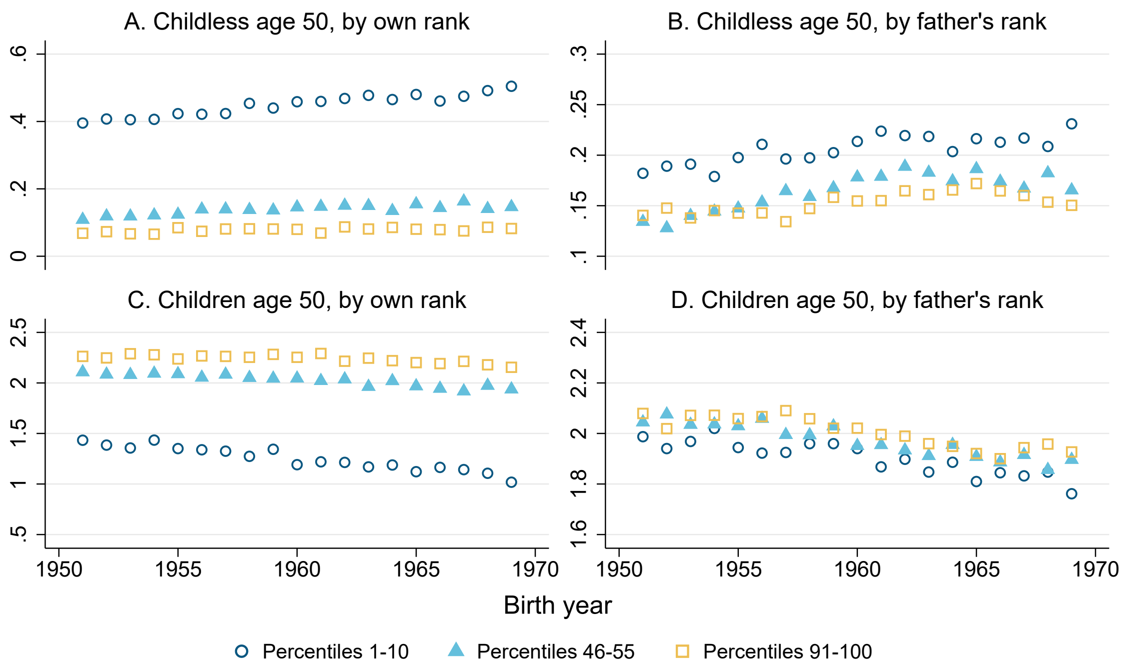

Notes: Scatter points represent ten percent of each cohort of Norwegian men born between 1951 and 1969.

Figure A.2: Sibling mean comparisons before and after firm bankruptcies, sampling at t-3.

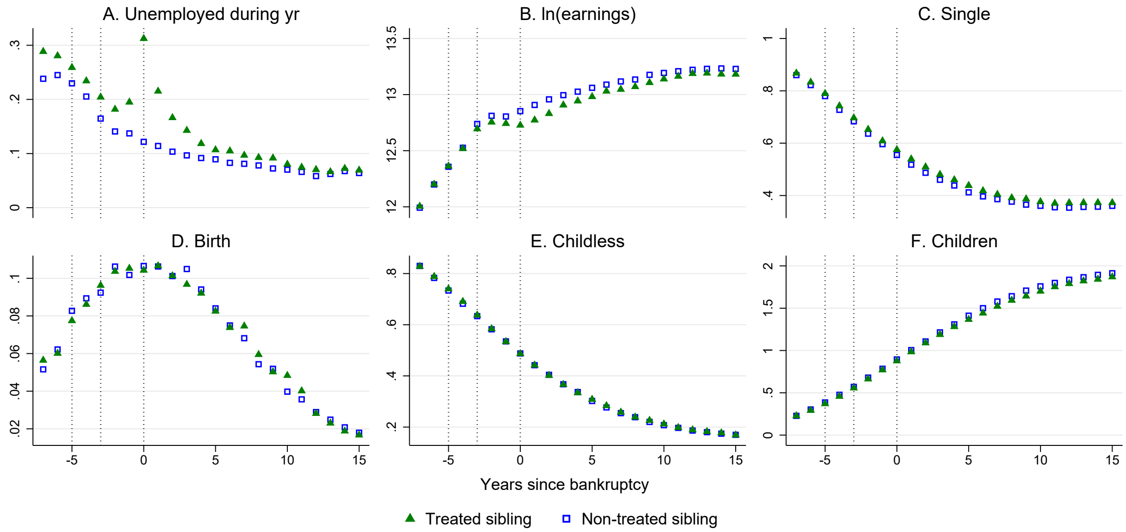

Notes: Vertical lines indicate year of observed November job (year -3), year of event (year 0), and reference year (-5). See text and notes to 14 for a description of samples.

Figure A.3: Sibling mean comparisons before and after establishment shutdowns.

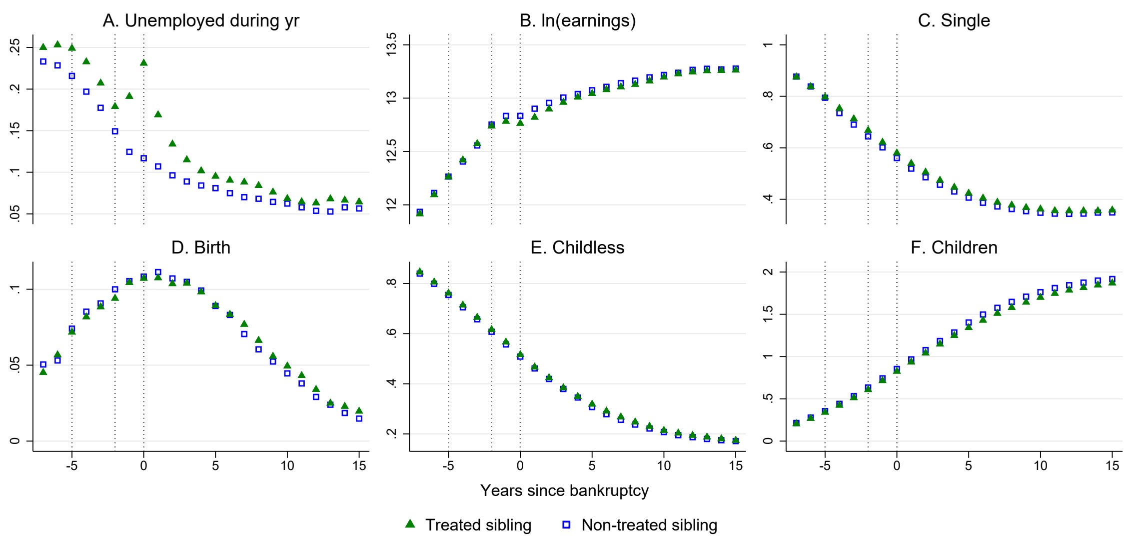

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). See text and notes to Figure 15 for a description of samples.

Figure A.4: Sibling mean comparisons before and after bankruptcies, balanced sample.

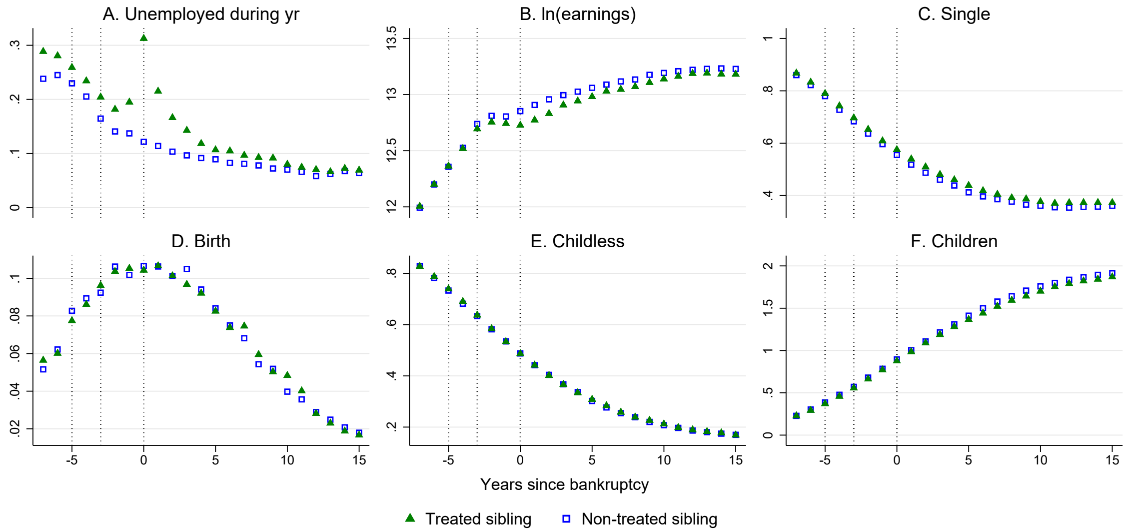

Notes: Vertical lines indicate year of observed November job (year -2), year of event (year 0), and reference year (-5). Sample of treated siblings consists of Norwegian-born men who in year -2 worked at an establishment that shut down two years later (between 1995 and 2004) and were age 25-35 in the 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 112 397 in the treatment group and 121 051 in the control group.

Figure A.5: Average education and unemployment over time

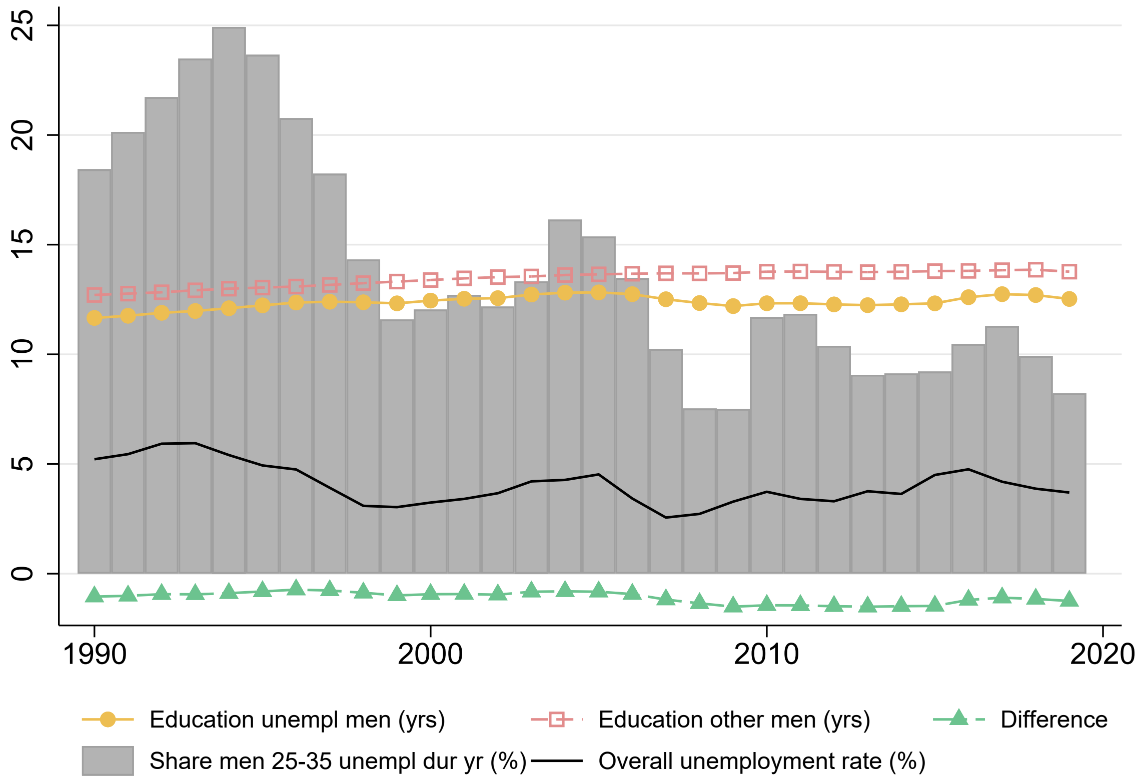

Notes: Vertical bars show the fraction of men aged 25-35 with registered unemployment during the prior year; and solid line shows the annual unemployment rate collected from https://data.oecd.org/unemp/unemployment-rate.htm. Scatter points depict mean years of schooling for those with and without registered unemployment, as well as the difference in attainment. Observation count is 8 881 215.

Figure A.6: The relationship between education and fertility over time, sensitivity to controls

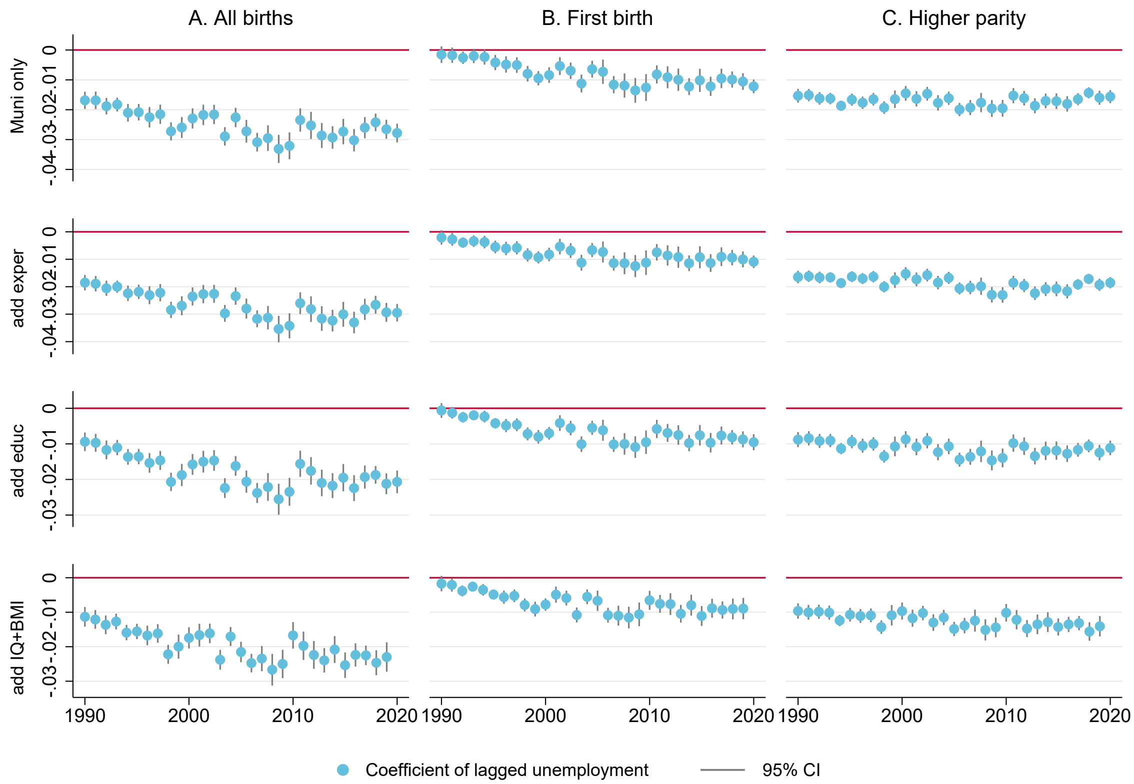

Notes: Scatter points show the estimated coefficient of individual unemployment status from a regression of birth on registered unemployment the prior year. All regressions control for year of observation. Regression in top row includes 428 municipality fixed effects; the second row adds polynomial of 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 1990-2019. Observation count is 8 881 215 (7 897 334 in bottom panels because of missing IQ or BMI data). Mean birth rate is 0.098 and mean registered unemployment is 0.143.

Figure A.7: Fertility across the earnings distribution.

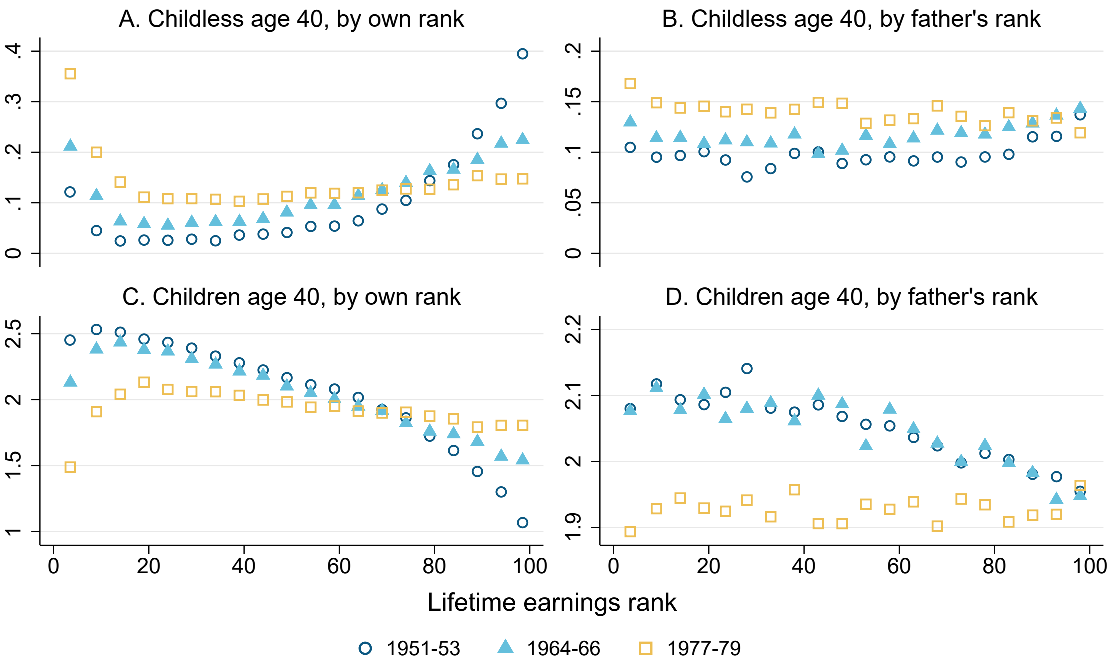

Notes: Each scatter point represents five percent of Norwegian women born between 1951-1953, 1964-1966, and 1977-1979, respectively. Observation count is 234 454.

A.2 Fertility and earnings distribution for women

Figure A.8: Inequality in fertility over time.

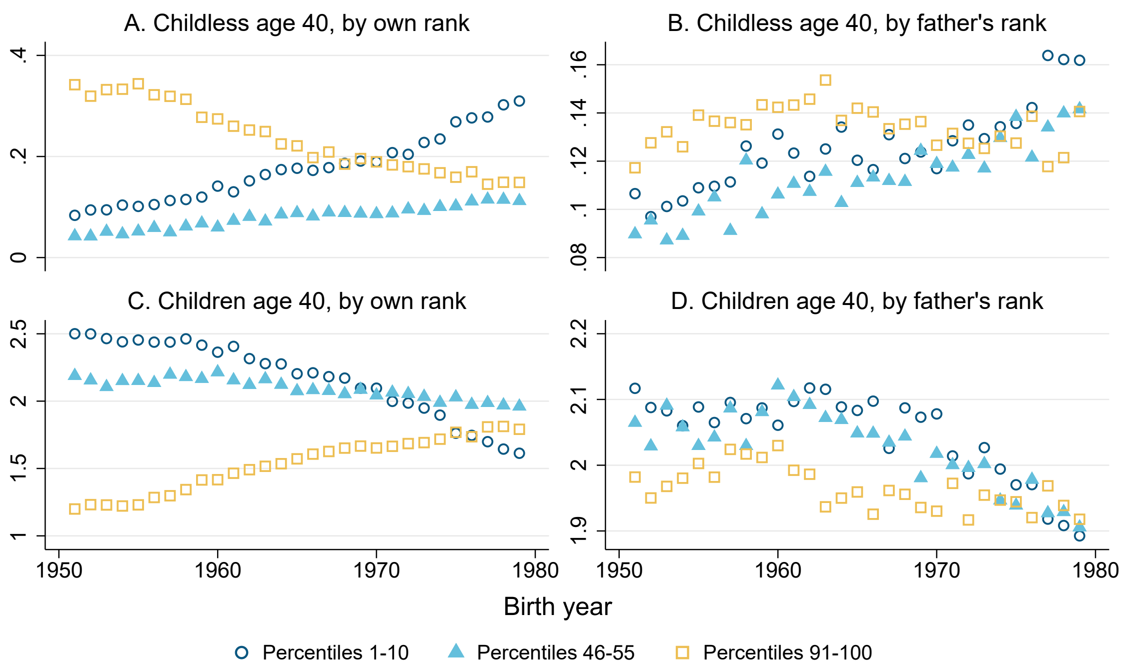

Notes: Scatter points represent ten percent of each cohort of Norwegian women born between 1951 and 1979.

Figure A.9: Absolute earnings over time.

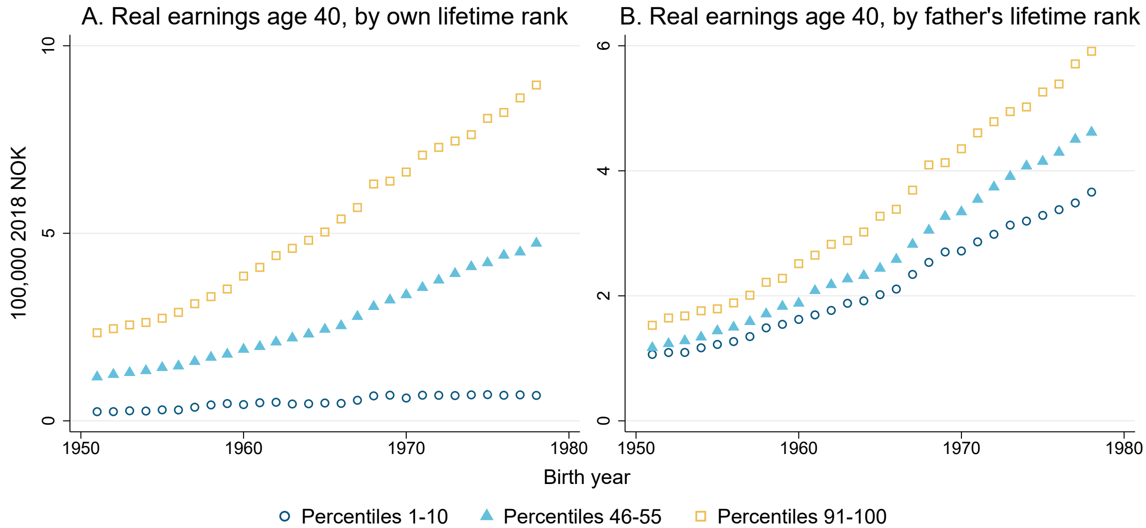

Notes: Scatter points represent ten percent of each cohort of Norwegian women born between 1951 and 1978. Earnings are observed at age 40, are inflated to 2019 NOK, and are depicted in units of 100 000. Observation count is 377 233.