Losing in a Boom: Long-term Consequences of a Local Economic Shock for Female Labour Market Outcomes*

Patrick Bennett† Chiara Ravetti‡ Po Yin Wong§

October 2021

Abstract

This article examines the long-term labour market consequences of a positive and large-scale economic shock, the discovery of oil and gas in Norway. Using longitudinal data on the entire Norwegian population, we find that the shock increases male income by around 7%, while reducing female income by up to 14%. Although married women experience the largest income losses, they also have higher household income, revealing the importance of labour supply adjustments within households. While these income shifts persist for two decades, the subsequent generation of female workers are able to close the income gap with their peers in areas less affected by the oil boom.

Keywords: female labour supply; natural resources booms; gender
JEL Classification: J1, J12, Q33

*The views expressed in this paper are those of the authors and do not necessarily reflect those of the Hong Kong Monetary Authority. The authors wish to thank Aline Bütikofer, Katrine Løken, Kjell Salvanes, Rick Van der Ploeg, Michael Ross, Gerhard Toews, seminar and workshop participants from HEC Montréal, the Norwegian School of Economics, the Centre for International Environmental Studies, NAERE, OxCarre, CEA 2019, ESPE 2019, EALE 2019, SIEPI 2020 and IAERE 2020 for their helpful comments and suggestions. This work was partially supported by the Research Council of Norway through its Centres of Excellence Scheme, FAIR project No. 262675. The authors acknowledge financial support from the Norwegian School of Economics, FAIR, the Research Council of Norway, Politecnico di Torino, Oxford Centre for the Analysis of Resource-Rich Economies (OxCarre), and the Hong Kong Monetary Authority. The paper was previously circulated under the title "Women and Resource Windfalls: The Norwegian Oil Boom".

†University of Liverpool and FAIR-CELE
‡Department of Management and Production Engineering, Politecnico di Torino
§Research Department, Hong Kong Monetary Authority

1 Introduction

Over the past few decades, the gender wage gap in industrialized countries has been narrowing and female labour force participation has been steadily increasing. However, this overall trend varies significantly across time and geography (Goldin, 2006; Olivetti and Petrongolo, 2016). Understanding the different labour market outcomes of men and women during economic transformations has important implications for gender equality, yet it is empirically challenging to study this question (Benshaul-Tolonen and Baum, 2019). Ideally, it requires a sizeable and plausibly exogenous change in economic conditions as well as longitudinal data on labour market outcomes over a long time horizon.

This article addresses this challenge by exploiting a substantial and persistent sectoral shock, Norway's first discovery of oil in 1969, to analyse the effects of an economic boom on individual labour market outcomes for men and women. We leverage the unexpected nature of the shock and the spatial variation in the development of oil-related industries within Norway to identify changes in individual income, household income, and labour supply. We also examine key mechanisms behind these changes, taking advantage of detailed individual- and household-level data that allows us to follow Norwegian residents over their life cycle, before and after the local shock. We hypothesize that the effects of the economic boom on the labour market depend on the substitutability of men and women in the sectors involved and on the reallocation of joint labour supply within households. In addition, the short run impacts may not persist in the long run, for instance, due to the greater flexibility to modify lifetime decisions such as full versus part-time employment, education, fertility, and occupational choices. We test these adjustment channels for different cohorts of the local population.

Following the discovery of oil and the subsequent boom, the income trajectories of men and women are starkly different. Our baseline estimates from an event study show that young working-age men earn significantly higher annual income in booming markets, while young working women, surprisingly, experience a significant fall in annual income. These changes persist for the following 20 years. Importantly, prior to the discovery of oil, differences in income across our treated and control regions are small in magnitude and statistically insignificant, indicating that there were no diverging income trajectories pre-boom. This result is robust to a battery of tests with alternative specifications of the spatial treatment units, to different definitions of the dependent variable, and is not driven by migration across different parts of Norway.

The fact that the income of women declines at times of great economic prosperity is relevant not only for gender equality broadly, but also for intra-household equality and decision-making. Since a large fraction of men and women make labour supply decisions jointly with their partner, we examine our core result not only for individual workers, but also in the context of households. We exploit data on married couples to show that intra-household dynamics are an important factor in explaining the observed decline in income for women. We find that women who are married before the oil shock experience the largest decline in individual income. However, their total household income increases as their husbands earn more. Overall, women's relative share of household income declines by 6%. It is difficult to evaluate the welfare effects of such a loss because the reallocation of tasks within the household and the higher aggregate income could benefit both spouses. However, women's lower share of household income is potentially problematic if it reduces their bargaining power (Basu, 2006), if it influences marital satisfaction and stability (Bertrand et al., 2015), and if gender inequality or female unemployment are linked to domestic violence (Aizer, 2010; Anderberg et al., 2016; Cools and Kotsadam, 2017).

We isolate the key mechanisms that drive the reduction in female income in booming regions. In particular, we show that the local economic shock decreases female labour force participation by almost 30%. This decline is driven by a reduction in full-time employment, while for men there is little adjustment at the intensive margin among those who are already attached to the labour force. We further show that men move into high-paying occupations following the boom, while women do not. Thus, the widening gender gap in income is driven mostly by women dropping out of the labour force or not having access to high-paid jobs, either directly in the oil sector or in related industries benefiting from its economic spillovers. Importantly, women's income losses do not seem to depend on their educational choices nor the decision to have children.

Finally, we extend the analysis to the next generation of workers, who were children at the time of the oil discovery and join the labour force only after the shock. These children are exposed to the economic shock in multiple ways at young ages, for instance through a rise in family income (Løken, 2010; Løken et al., 2012). We show that the income losses for women progressively disappears for the following cohorts of female workers. The youngest cohort of girls close the pre-existing lifetime income gap with their peers in control regions. In contrast, male lifetime income remains systematically higher for early and later cohorts. These results complement our main findings with an important long-term outlook. Among men and women already in the labour market, the effects of the oil boom are long lasting. Among the following generation of workers, the economic gains persist for men, while the income losses are only present among the first generation of working women.

Our study contributes to a growing literature on the spillovers from localized industrial activities that uses within-country studies (Greenstone et al., 2010). In particular, a number of articles exploit the fact that natural resources offer a source of exogenous variation in economic conditions that are associated with increases in local economic activity and employment during booms.[1] Few studies, however, adopt a within-country design to compare local labour outcomes for female and male workers, with evidence on inequalities among workers and within households even scarcer. One exception is Aragón et al. (2018), who find that coal mine closures crowded out female employment in manufacturing and services at the district-level, while male employment rose, with persistent effects over two decades.[2] Our paper expands this prior literature by emphasizing the differential effects of a local economic shock on male and female workers at the individual level and even inside households.

Our study also provides new evidence on the long run effects of a resource-driven boom on individual labour market outcomes of men and women. Existing studies, such as the recent literature on fracking in the US, have leveraged an analogous economic boom, but only explore short to medium-run changes in the labour market. These studies find an overall positive impact on wages and other labour market outcomes (Allcott and Keniston, 2018; Bartik et al., 2019).[3] A notable exception is the work by Jacobsen and Parker (2014), which uses US county-level data from the 1970s and 1980s to show that oil booms have long-term negative consequences in terms of joblessness and income depression when a boomtown experiences a bust. For Norway, a small body of work examines the long run consequences of the oil discovery on individuals, but gender and within-household decision making are rarely considered.[4]

Another contribution of our study is the novel finding that women's annual income declines

[1]See for instance Black et al. (2005); Michaels (2011); Marchand (2012); Fleming and Measham (2015); Autor et al. (2019); Bartik et al. (2019). One key issue examined by this literature is whether the resource sector offers some complementarities with other sectors, or instead crowds out manufacturing and other labour-intensive industries.

[2]In addition, Maurer and Potlogea (2020) show that oil discoveries in the US in the early 1900s increases labour force participation among single women, while Kotsadam and Tolonen (2016) show that during mining booms in Sub Saharan Africa, women shift away from the agricultural sector into services, but overall female employment declines.

[3]These studies also find significant—but not necessarily positive—social changes, such as increases in fertility (Kearney and Wilson, 2018), more high school dropouts (Cascio and Narayan, 2015), increased migration (Wilson, 2020), and higher crime rates (James and Smith, 2017; Bartik et al., 2019).

[4]Løken (2010) and Løken et al. (2012) examine the relationship between family income following the oil boom and children's educational outcomes, while Bütikofer et al. (2018) assess how intergenerational mobility among men changes following the discovery of oil.

and remains depressed following the boom while men's income soars. This finding is especially surprising since Norway is considered to be one of the most egalitarian countries in the world in terms of gender norms (WEF, 2020; UNDP, 2019). One of the largest shocks studied in this literature is World War II: following manpower mobilization towards one of the most gendered industries, the military, regions with higher mobilization rates experience an increase in female labour supply (Acemoglu et al., 2004).[5] The effects, however, are long-lasting only for highly educated women who could enter white-collar occupations (Goldin and Olivetti, 2013). The spillovers from these changes in labour demand are relevant also for the following generations of women (Fernández et al., 2004) and are linked to lower educational attainment and more precocious marriage and fertility choices (Jaworski, 2014). Given that the Norwegian oil boom is a substantial and long-lasting sectoral shock with general equilibrium effects on other sectors, our findings highlight the importance of the dynamics of women's income following a positive shock in a male-dominated industry.

Finally, our study contributes to the diverse literature on the responsiveness of spousal labour supply in compensating for labour market shocks and the importance of the "added worker effect" (Lundberg, 1985). While most papers find evidence of a household labour supply response following spousal earnings shocks (Stephens, 2002; Goux et al., 2014; Blundell et al., 2016; Autor et al., 2019), others find a limited role for spousal adjustments to smooth adverse shocks (Halla et al., 2019). In line with the majority of this literature, our findings point to a significant role for within-household labour supply adjustments in response to local economic shocks. While the Norwegian oil boom occurred at a time of increasing female labour force participation, our findings are informative for understanding gender equality in the presence of economic shocks.

[5] The consequences of negative labour market shocks for women have also been documented. For instance, a study of Austrian plant closures finds that job displacement reduces income and job-specific human capital for women in highly skilled occupations, also decreasing average fertility by 5-10% (Del Bono et al., 2012).

The rest of the article is organized as follows. Section 2 illustrates the analytical framework we use to formulate our testable hypotheses. Section 3 discusses the Norwegian oil discovery. Section 4 presents the data and the estimation sample, and Section 5 outlines the identification and estimation strategy. Section 6 discusses the results and potential mechanisms. Section 7 concludes.

2 Analytical Framework

This section discusses the analysis of labour market outcomes for different types of workers in the aftermath of a resource boom within classic frameworks of structural change and Dutch disease la Corden and Neary (1982). The sudden discovery and development of a sector like oil and gas is expected to cause significant shifts in the labour market. First of all, a sectoral boom causes an expansion in the previously non-existent oil industry and in related activities, such as construction work to build the extractive infrastructure. Therefore, local labour demand in these sectors will rise. However, male and female workers are imperfect substitutes in these productive activities and the oil industry's labour demand mostly targets the male workforce, both due to their individual characteristics (such as physical strength) and because of cultural norms.[6] Moreover, the increase in local labour demand could also attract foreign workers, so that the booming regions might experience population growth.

Second, the resource boom can also generate economic spillovers towards other industries, particularly non-tradable sectors like local services (Black et al., 2005; Marchand, 2012; Jacobsen and Parker, 2014). An expansion of these "spillover sectors" plausibly creates an increase in labour demand for both men and women, countering the effect of the male-dominated oil

[6] Workers in North Sea oil and gas installations typically face demanding working conditions that make offshore jobs more inconvenient for women: periods away on a platform typically last around 2 weeks, in cramped spaces with limited privacy, and entail physically demanding shifts of 12 hours (Parkes, 2012).

industry. Women could even have an advantage in these spillovers sectors: empirical evidence has long established that a growing service industry increases labour opportunities for women, and can substantially narrow gender gaps in working hours and wages (Ngai and Petrongolo, 2017). Lastly, it is also possible that an appreciation of the exchange rate following the exports from the extractive industry would crowd out other exporting activities and destroy some jobs, but in most cases there is limited evidence for these kinds of negative effects, and actually these sectors may benefit from some agglomeration spillovers (Michaels, 2011; Allcott and Keniston, 2018).

The overall effect on male and female labour outcomes depends on these relative sectoral effects, as well as the responsiveness of men and women's labour supply. The adjustments in the labour supply depend not only on individual workers' choices, but also on intra-household decision making regarding time allocation. One of the key contributions of our analysis is precisely to examine the dynamics of the labour supply of men and women, both in the short and long run. We account not only for the asymmetric labour demand shock across sectors, but also for the labour supply response of individuals and households. Overall, we test three hypotheses.

1. Male wages increase

In the short run, given that men in the pre-boom equilibrium are almost already at full employment,[7] the higher demand for male labour from the oil industry should lead to higher average wages for men, in order to attract workers already employed in other occupations. The equilibrium effect on their aggregate earnings depends on the elasticity of the labour supply, namely on how many more hours men chose to work and if migrant workers were available to fill in the new jobs.

[7] In 1970, 92% of all men born 1940–1949 are employed with some hours of work.

2. Female labour force participation and wages may increase or decrease

The effect of the oil boom on women is instead theoretically ambiguous. On the one hand, even if women do not enter oil-related industries, the non-tradable sectors that benefit from local spillovers could absorb female workers, increasing their employment rates and offering better wages than in low oil regions. On the other hand, for the majority of households, women's labour supply is typically determined within couples as a residual, since the husband's labour supply is rigidly fixed at full-time employment, while the wife's labour supply adjusts depending on family's preferences, conditions and needs (Donni, 2007). Therefore, married women might experience a substitution effect from intra-household decision processes, because after the economic boom they could stay home and work in household production (housekeeping, child-care, and so on) as their partners start earning higher incomes. Ultimately, it is an empirical question whether the positive effects from greater opportunities in the non-tradable sectors and services prevail over intra-household substitution effects.

3. Short run results may not persist in the long run

In the long run, we expect greater substitutability between men and women both on the labour demand and on the labour supply side: cultural norms could change, workers of both sexes could decide to acquire different forms of training and education, career and family planning could be adjusted more freely. Therefore, we do not expect the short run results to necessarily apply in the long run. We hypothesize that, given a higher elasticity of substitution between workers in the long run, young generations of men and women joining the labour force after the first economic boom should not present as much heterogeneity in labour market outcomes as the first generation of workers impacted by the sectoral disruption.

For women in particular, more economic opportunities could appear if, over time, cultural norms evolve in households—the supply side— and oil industries—the labour demand side

(Fernández, 2013). Substitution effects inside the household, however, can theoretically evolve in different directions over time: it could become standard in oil regions for women to stay at home and for men to work in oil-related activities, or it could be that the extra wealth brought to the household by the economic boom catalyses more transformations, such as education opportunities for young girls that later induce better labour market outcomes for women. Overall, it is an empirical question to establish which of these effects dominates in the long run.

3 The Discovery of Oil and Gas in Norway

The discovery of substantial offshore oil and gas deposits[8] in 1969 was an unexpected event in the history of Norway and thus offers an attractive setting for studying how the impacts of a large economic shock on labour market outcomes differ by gender. Up to the 1960s, the Norwegian economy relied mostly on fishing, shipping, forestry and agriculture. In those years, the prospects to develop an oil industry seemed limited - for instance, in 1958, the Geological Survey of Norway wrote to the Ministry of Foreign Affairs that the possibility of finding oil off the Norwegian coast was to be discounted. In 1966, the first exploration well proved to be dry. In 1967, the first actual oil discovery on the Norwegian shelf was not economically viable and it took another 30 years before the field was ultimately developed (Norwegian Petroleum, 2020). Then on December 23, 1969, the American oil company Phillips Petroleum announced to the Norwegian government the discovery of Ekofisk, the first 'giant' oilfield in Western Europe (Van Den Bark and Thomas, 1981).[9] This event marked the beginning of the Norwegian oil industry, which substantially changed the local economy in the following decades, with long-

[8]For brevity, throughout the article we refer just to oil, but the resource discovery and the industry that developed around it also comprises a substantial natural gas sector.

[9]An oilfield is defined by the American Association of Petroleum Geologists as 'giant' with at least 500 million barrels of oil equivalent in ultimate recoverable reserves before extraction starts, including condensate and natural gas (Horn, 2011).

lasting socio-economic consequences (Bütikofer et al., 2018).

The discovery of any giant oil field constitutes a significant economic discontinuity, with major macroeconomic repercussions (Arezki et al., 2017). Oil findings of these dimensions are rare and unpredictable: despite significant investments in exploration, finding a giant oil field is rare and plausibly exogenous (Lei and Michaels, 2014).[10] For Norway, the first giant discovery in 1969 was even more unexpected since exploration companies had been drilling unsuccessfully in the Norwegian continental shelf for more than five years. Phillips Petroleum was the last company still exploring in 1969 and it seems that it would have abandoned its last exploration, had it not been cheaper to drill than to pay a fine to the Norwegian oil office for abandoning its exploration commitments (Sandbu, 2009). We therefore consider the discovery of the first giant oilfield as a random economic shock, exogenous in its timing, and impossible to anticipate by the local population. A series of major oil discoveries was made in the following decades, but at that point it had already become clear that Norway could develop its own oil industry. Thus, our empirical strategy relies only on the initial offshore discovery of oil in 1969 as an unanticipated event, but not any of the subsequent oil discoveries.

The sudden oil boom initially affected only the few industries and areas specializing in the extraction and refinery of oil and gas. Following the first find, the Norwegian authorities designated as the country's main oil base Stavanger, the largest city in the Rogaland region located relatively near the offshore Ekofisk discovery. Rapidly, many other sectors such as construction and certain types of manufacturing also benefited from a local economic boom following the discovery of oil. While the 'Dutch disease' could have destroyed some manufacturing industries, Norway managed its petroleum sector carefully so to avoid the typical

[10]To see how unpredictable a giant finding is, consider that, since 1965, only 2% of all wells drilled around the world resulted in giant oil or gas discoveries (Toews and Vezina, 2017). One exemplary illustration comes from the Johan Sverdrup oilfield, one of Norway's largest discoveries, which was found in 2010 few metres away from an exploration well where the French company Elf Aquitaine drilled in 1971, without finding any oil (Kavanagh, 2013).

negative consequences of resource discoveries (Larsen, 2006). Moreover, Norway did not have a sizeable manufacturing export sector before the oil boom, and mostly relied on primary goods production like fish and timber, so the crowding-out of other local manufacturing activities was not substantial (Bütikofer et al., 2018).

Figure S1 in Appendix A plots the location of the Ekofisk oil field relative to the city of Stavanger. Ekofisk is an offshore oilfield located almost midway between the UK and Norway. The distance from Stavanger is about 320 kilometres, and the closest location on Norwegian land is 260 km away from the oilfield. Statoil, Norway's state-owned energy company, and the Norwegian Petroleum Directorate were established with their headquarters in Stavanger in 1972, soon followed by other companies (Løken, 2010). Over time, the spillovers from the industry spread beyond the first oil location. For this reason, we discuss carefully the definition of 'oil regions' in Norway in section 5.1.

Even if the offshore oil and gas industry is not labour intensive, employing just above 80,000 workers in the whole country (2.5% of the working age population) at its peak in 2014,[11] the two municipalities with the highest ratio of employees in the petroleum industries, Sola and Stavanger (both in Rogaland), had 15.7% and 13.6% of local employment in oil-related industries, respectively (Ekeland, 2017). Therefore, the local shock should have generated some significant labour market and income effects, possibly with spillovers onto other local activities, such as manufacturing and services.

4 Data

[11] Authors' calculation from https://fred.stlouisfed.org/series/LFWA64TTNOA647S.

4.1 Individual Data

To examine the impacts of the local demand shock from the oil discovery on individual outcomes in the short and long run, we use annual register-based data as well as decadal Census data. Each of these data sources are described below.

4.1.1 Registry Data

Our main data source is the Norwegian Registry from Statistics Norway, which covers the entire population and follows individuals annually from 1967–2017. Using an anonymous personal identification number for all Norwegian residents, we merge data on individuals across three different registers from different administrative sources to create an individual-level panel. Importantly, individuals remain in the data at all points in time irrespective of whether they are employed or not.

First, we observe individuals in the central population register (*folkeregisteret*) yearly starting in 1970. The population register contains data on an individual's municipality of residence, gender, birth year, and marital status, including date of marriage. Prior to 1970, yearly snapshots are available in 1965 and 1967 with the same information.

Second, from the tax and earnings register, we observe annual income, one of our key outcome variables, every year dating back to 1967, prior to the first discovery of oil at the end of 1969. This measure of income comprises total labour income, including any income earned from self-employment, as well as any taxable benefits received during the year.[12] We cannot extrapolate actual wages to construct a gender wage-gap, so throughout the article we refer to annual income as defined above. Income throughout is measured in 1998 Norwegian Kroner. Third, using the education register, we can define an individual's education as the

[12]Taxable benefits include public transfers such as parental leave, unemployment, or sickness benefits. These benefits should not significantly affect our analysis, since in Norway they are determined at the national level, and thus do not differ systematically in relation to treatment status post-oil.

highest completed level of schooling.

4.1.2 Census Data

Our analysis also relies on data from the 1960, 1970 and 1980 Census, which provide detailed information on individual employment, such as occupations classified according to the International Classification of Occupation (ISCO), industry of employment classified according to the International Standard Industrial Classification of All Economic Activities (ISIC), employment status including self-employment, as well as annual hours worked. Data on economic activity is collected based on the 12 months preceding the Census date: for example, hours worked in the 1980 Census is an annual measure of hours from November 1979 to November 1980. The Census has near universal coverage, including all individuals registered as residing in Norway at the start of November in the Census year. While the yearly panel from the Registry is essential to study the annual evolution of individual income after the oil shock as well as to inform the similarity of those in treated and control areas prior to the discovery of oil, the Census data allows us to explore other outcomes, such as labour market attachment and occupational changes. In addition, with data going back to 1960, the Census data permits the examination of pre-discovery trends over a longer time horizon compared to the yearly panel.

4.2 Sample

The first part of our analysis focuses on a sample of young individuals born between 1940-1949 (aged 20-29 at the time of the local boom in 1969). We focus on these birth cohorts as they are already active in the labour force when the oil and gas industry starts to develop, but are still in an early career phase and could make significant changes to their work choices. By excluding younger individuals, we minimize the influence of the oil boom on decisions regarding human capital investments versus joining the labour force. By excluding older individuals, we abstract from the effects of the oil boom on retirement decisions when we look at long-term labour outcomes. Moreover, key family decisions (such as household formation and fertility) are set relatively early in the life cycle: for instance, the median age of first childbirth is 22 for women born between 1940–1949. As a first baseline, we restrict the sample to those workers who have strong attachment to the labour force prior to the oil shock, namely individuals who have a positive income in all three years over the period 1967–1969.[13]

We also consider individuals with different working statuses before the boom, including those who are not employed at all, to examine all possible changes in labour force participation along the extensive and intensive margins. While our main focus is on the sample of workers attached to the labour force, results throughout the article are unchanged if we consider instead all individuals irrespective of labour force attachment. In the last part of the empirical analysis, we examine labour outcomes for younger cohorts that are exposed to the oil shock during childhood or adolescence and start working afterwards.

5 Empirical Strategy

We now present our empirical approach to analyse the effects of the oil boom on employment in Norway. First, we discuss the identification strategy and how we approach any threats to identifying causal effects. Afterwards we outline our empirical model.

5.1 Identification - Defining Booming Labour Markets

To evaluate the effects of the Norwegian oil boom, we divide Norway into different areas that we assign to treatment or control groups depending on the spatial variation in the importance of

[13] About 34% of women have positive income in all three years before the oil shock.

the oil sector and related industries. Our unit of analysis is a 'local labour market', a geographic segmentation that organizes Norway into 46 different units (Bhuller, 2009), constructed on the basis of commuting patterns between individuals' municipality of residence and work. These markets are defined irrespective of administrative boundaries, so they are less affected by political decisions than counties and municipalities. In fact, municipalities in the same local labour market may belong to different counties.[14] This way, our assignment of geographic areas to treatment and control groups is less susceptible to political changes that could relate to oil and administrative decisions made within each county.

To measure the intensity of exposure to the oil shock at the local labour market level, we move away from a simple measure of spatial distance from the location of the oil discovery, because the Ekofisk oil field was located offshore. Offshore oilfields are not expected to have the same local effects as onshore wells (as confirmed by the empirical literature, see for instance Cavalcanti et al., 2019). The nearest coastal area is not necessarily the most impacted location: other places might also benefit, conditional on the ease of transporting the hydrocarbons to a near harbour and then in moving them through the appropriate infrastructures to refineries or other industrial locations.

For our core identification strategy, we follow Bütikofer et al. (2018) and use the share of employment in oil and oil supply industries in 1980 as our measure of exposure to the oil shock. Since there was virtually zero oil production in 1970, the year 1980 represents the first available period after the initial discovery in which data on the oil industry is available. We include industries directly involved in the extraction and refinery of petroleum, as well as secondary industries that support and supply to oil production, namely crude petroleum and natural gas production, petroleum refining, manufacture of products of petroleum and coal, manufacture

[14]There are approximately 420 municipalities in Norway, and the average local labour market contains 9 municipalities.

of machinery (including oil rigs), manufacture of transport equipment, and construction other than building construction (including oil well drilling). Brunstad and Dyrstad (1997) show that supply industries should be considered as part of the oil sector since they also face a positive demand shock from the oil boom in the 1970s.[15]

Figure 1 depicts the spatial distribution of the intensity of the oil treatment. As in Bütikofer et al. (2018), we define three levels of exposure to oil: high, medium, and low. We define a local labour market to have high exposure to oil if more than 10% of its employment is in the oil sector, corresponding to the top quartile of oil employment across all labour markets. The average employment share in oil among all high oil regions is 14% and the two regions with the highest shares are Kongsberg, which had 26% of employment in oil, and Sunnhordland, which had 22% of employment in oil. Medium exposure is defined as those with between 7.5% and 10% employment in oil and low exposure is defined as having less than 7.5% of employment in oil.

In the analysis of individuals and households, we assign each person to a local labour market based on residence in 1965, avoiding potential biases from endogenous relocations. Table 1 displays the characteristics of our initial sample of young workers by different types of labour markets. All values are measured in 1969, with males in the first column and females in the second column. In the sample of young workers, 15–16% live in a high oil local labour market in 1965, 19% in a medium oil labour market, and the vast majority in a low oil one. There is a stark difference in the average income levels between genders, with men earning about 50% more than women. Additionally, women participate less in formal labour markets prior to the oil discovery, as the sample of men is substantially larger.

Table S1 reports a detailed breakdown by gender of the employment shares of the industries which comprise the oil sector using the estimation sample from the 1980 Census as described

[15]This definition of oil sector employment is similar to the one suggested by Allcott and Keniston (2018).

in Section 4.2. While the production and refinery of petroleum is an important industry for men in high oil areas, the manufacture of oil rigs and oil well drilling comprises a higher share of total employment. On the other hand, substantially fewer women in both high and low oil areas tend to work in the oil sector compared to men.

In sum, our identification relies on (i) the exogenous discovery of oil in 1969 as our treatment, (ii) the expansion of oil industries resulting from this shock in each local labour market as our source of spatial variation for treated and control areas, and (iii) on the assignment of individuals to a local labour market based on their residence before the oil shock.

5.1.1 Threats to Identification and Alternative Strategies

Our main identification strategy compares high oil (treated) local labour markets to low oil (control) ones using the industry-based measure of exposure to oil. Such a comparison over time relies on the parallel trends assumption—that in the absence of the oil discovery, high and low oil local labour markets would have evolved similarly in terms of labour market outcomes. While an industry-based measure best encompasses all industries which expanded as a result of the oil discovery, there could be two potential—and interrelated—concerns with our definition.

First, while the parallel trends assumption does not require that variation in high/low employment shares in oil across local labour markets be unrelated to pre-existing area-level factors, an industry-based measure may reflect a comparison between regions which had specific pre-existing characteristics, which are exacerbated by the discovery of oil. The correlation of oil developments with pre-existing characteristics of local labour markets is a concern only if it were also linked with how the gender gap evolves, but nonetheless it is informative to understand the underlying reasons behind why employment in particular areas becomes dominated by the oil industries.

To address this issue, we show that individuals in high and low oil regions did not have systematically different trends in the outcomes of interest prior to the oil boom. Factors such as female and male labour force participation evolve similarly from 1960 to 1970 between high and low-oil local labour markets (Figure 2). The lack of differential pre-trends confirms that, in the time leading up to the discovery of oil, trends between high and low oil areas are similar. Moreover, we document that area-specific characteristics such as industry structure, population density, labour force participation, gender composition, and education levels measured in 1960 prior to the oil discovery explain only to a limited extent the development of oil industries. For instance, while high oil labour markets are concentrated near the first oil discovery and on the coast, large urban centers are equally dispersed across the different levels of oil employment shares: Oslo (the capital) is low oil; Bergen (the second largest city on the west-coast) is middle oil; and Stavanger (in south-west Norway) is high oil. Indeed, Tables S2 and S3 show that area-specific characteristics in 1960 have little predictive power and are insignificant in determining high oil status.[16]

Second, the classification of local labour markets into high, middle and low oil regions is based on employment shares measured in 1980, thus incorporating all development in the oil industry that occurred after the initial discovery. While the random nature of the oil discovery is reflected in the fact that pre-existing differences fail to explain future high oil areas in 1980, an alternative approach in measuring exposure to oil would exploit the direct spatial component of the discovery, relying solely on ex-ante factors. However, we choose to adopt an industry-based measure rather than a spatial measure as our baseline because, unlike the case of onshore

[16]Results reveal that pre-existing area-level factors in 1960 such as the most common 2-digit industry shares, broad industry shares, population density, and educational levels have little predictive power for high oil status in 1980 (Appendix D). An exception is employment in manufacturing of metal products, machineries, and equipment, which is positively related to high oil in 1980. This 2-digit industry includes shipbuilding, which directly benefited from the discovery of offshore oil. Even so, industry can only predict 12% of the likelihood of having a high share of employment in oil in 1980, and our results are also robust to excluding local labour markets which specialize in such industries.

shale plays in fracking, distance is not necessarily the most important predictor for the level of extractive activities that follow a discovery offshore. Furthermore, in our baseline identification strategy, we include local labour market fixed effects, which capture among other things distance to the offshore deposit. As a robustness test, we also perform a direct spatial comparison of the closest local labour markets to the first offshore field with the rest of Norway, relying on the 'as good as random' nature of the location of the first oil discovery.

For completeness, we estimate a variety of alternative specifications to establish the robustness of our baseline results. Importantly, we rule out geographic mobility as a driving force behind our results. Theoretically, large movements of the population following the labour shock could distort our results: for example, the literature from the USA documents significant inflows of male workers in fracking areas (Wilson, 2020). Since we allocate individuals on the basis of their *initial* residence, we may understate the male gains from oil, since men that we consider in the control group could have migrated to high-oil to make more money. For women, however, the direction for any possible migration bias is theoretically less clear: they could be moving with men, or they could move away from a region offering only traditionally "masculine" jobs. In the analysis, the data allow us to identify movers and to show that results are unchanged when excluding them. These results, as well as a battery of other robustness checks regarding the capital Oslo and Northern regions, coastal areas, university and college towns, the most religious areas, and the presence of kindergartens, are presented in detail in section 6.2.

5.2 Empirical Model: The Impact of Oil Discovery on Income

The baseline empirical model for our analysis is an event study approach to examine the effects of the initial Ekofisk oil discovery in 1969 on the annual income of young cohorts (individuals born between 1940-1949). This approach allows the labour market impacts of the first oil discovery to vary over the years since the discovery of oil in 1969. Hence, we can trace the persistence of the shock over time. We estimate the following regression separately for men and women:

$$Y_{ibct} = \sum_{k=-3}^{10} \delta_k Oil_{ic}^k + \psi_c + \tau_b + \gamma_t + u_{ibct} \tag{1}$$

where Y_{ibct} denotes the log of annual income of individual i of birth cohort b who resides in local labour market c in 1965, measured at time t with $t = 1967, \ldots, 1980$. The index k represents the number of lags (forward or backward) from the time of the oil discovery in 1969 ($k = -1$), when our treatment starts by design. ψ_c captures time-invariant characteristics of each local labour market, including distance to the first offshore oil find and any other fixed characteristic of the market. τ_b represents birth cohort-specific characteristics, and γ_t denotes year fixed effects. The income level in the year of oil discovery is normalized to zero.

To examine the labour market effects of exposure to oil, we define the treatment Oil_c^k to correspond to residence in high oil local labour markets in 1965 and compare it to the control group of individuals residing in low oil areas. The estimated coefficients δ_k represent the average difference between individuals in high oil relative to low oil local labour markets in each year relative to the same difference prior to the oil discovery at $k = -1$. Pre-event coefficients δ_{-3}, δ_{-2} reveal any inherent differences in the evolution of income prior to oil discovery. We cluster the standard errors at the level of the initial (1965) local labour market of residence, assuming error independence across these markets.

6 Results

Figure 3 plots the evolution of the estimated δ_k coefficients in equation (1). Panel A presents results for males while Panel B presents results for females. As discussed in section 5.1, the validity of our approach rests on the parallel trends assumption: in the absence of the oil discovery, income dynamics in booming labour markets would remain parallel to those in labour markets that are not affected by the oil discovery, so that low oil markets provide a valid counterfactual for what would have happened in the absence of oil. The lack of significant differences in the income trends between individuals living in treated and untreated labour markets prior to the discovery of oil shows that pre-discovery income trajectories are similar, supporting such an assumption.

For the years after the oil discovery, the impact of the sectoral boom on the incomes of men and women is different, and these differences persist up to ten years after the discovery. For men in high oil regions compared to men in low oil regions, the boom takes a few years to translate into an increase in income: the estimated impact is not significantly different from zero (at the 5% significance level) for the first year after discovery. This is reasonable considering the lag between the discovery and the actual start of production in June 1971. From two years after the oil discovery, the estimated impact on income continues to increase and, ten years after the discovery of oil in 1980, the annual income for men in high oil regions is 6.4% higher than those in low oil regions, relative to the year prior to the discovery of oil ($t = 1969$).

A different pattern emerges for women. Similar to men, the oil shock also takes time to translate into changes in income for women, likely due to the time lag in oil production from the initial discovery. However, income is significantly *lower* for women in high oil local labour markets compared to low oil ones, from four years after the discovery and onwards.[17] The

[17]Our measure of income is not deflated using local prices, due to the lack of data on a regional CPI within Norway in the time period considered. It is plausible that prices in high oil markets become higher than those in

timing of this sharp fall in female income follows the steady increase in male income, suggesting that some intra-household reallocation of labour might be at play: as men are employed in high-paying oil jobs offshore or in other oil-related industries, their wives possibly cut back on labour force participation and accept lower annual income. We will test this mechanism in Section 6.1. Overall, compared to those residing in low-oil local labour markets, women in high oil regions earn significantly less after the boom. For the average young woman in a high oil region, income ten years after the oil discovery is 10% lower than that of the average young woman in a low oil region. At its lowest point, the income of women in high oil regions is almost 15% lower than in low-oil labour markets.

By taking the log of income as our dependent variable, we focus on the workers with positive annual income and who are thus participating in the labour force in a particular year. Results in Table S2 in Appendix C repeat the analysis with income measured in levels rather than logs, thus including all individuals in our baseline sample, even if they have zero income post-discovery. Measuring income in levels results in similar patterns for both men and women, though the decline in income among women occurs earlier when income is measured in levels. Moreover, Figure S3 confirms that women in high oil markets are less likely to have positive income after the discovery of oil. We examine the decision to work as a mechanism underlying our main result in Section 6.3.1.

6.1 Intra-Household Dynamics

As discussed in Section 2, the effects of a resource boom on different types of workers depend on the interaction between labour demand and labour supply. While this is true for all categories

low oil areas because of the economic boom, so that even if annual incomes of men in high oil local labour markets increase, as we document here, we cannot definitively conclude that their local purchasing power becomes higher as well. On the other hand, if prices grow more in high oil local labour markets, then our results on women's reductions in income become even more poignant.

of workers (skilled/unskilled, young/old, and so on), the fundamental issue in the case of gender divisions is that workers choose how to supply labour not only as individuals, but by and large, also jointly within couples. Following the boom, household income may be unchanged or even increased among couples if higher male earnings compensate for the declines in female income.

To inform such intra-household dynamics, we proceed in two steps. First, we separate the baseline sample of individuals in Figure 3 into married and unmarried individuals as of 1969. If the household is indeed a relevant unit of decision-making, then the income declines observed among married women should be even larger. Second, in order to observe the evolution of household income over the decade, we construct a sample of 'stable' couples, defined as those who are married continuously since 1969 (or earlier) and continue to remain married through to 1980. This subsample of 'stable' couples allows us to define a meaningful measure of household aggregate earnings, excluding the income of individuals who divorced during that time frame. Thus, we can compare long term changes in individual income to changes in household income and ask whether women's position within the household is impacted by the discovery of oil.[18]

Figure 4 compares the effect of the local shock on income among those who are married or unmarried at the time of the oil discovery. Our measure of marital status prior to the oil discovery abstracts from later marriages that follow the economic boom and hence are potentially endogenous. It is important to note, however, that the majority of men and women who are unmarried in 1969 eventually become married by 1980 (67% and 71% respectively). Bearing this in mind, both married and unmarried men benefit from the economic shock in terms of income, and the effects are virtually identical between the two groups (panel A).

For women, instead, marital status prior to the discovery of oil matters for the magnitude

[18]We obtain similar results after relaxing the stable couple restriction, and calculating household income irrespective of whether the household remains together by 1980, since divorces were not so frequent in those times.

of the income losses experienced. Prior to the oil boom, unmarried women earn approximately the same amount as married women in 1969. However, those who are already married in 1969 experience a sharper decline in income from the fourth year since the oil discovery relative to unmarried women. During the 4–9 year period after the discovery of oil, income losses among women who are already married in 1969 are double in size relative to those who are initially unmarried. Note that these single women are likely to marry over the course of the decade, and make the same intra-household allocation decisions as their already married counterparts, and hence also experience income decline, relative to those in low oil, albeit to a smaller degree. Overall, these results conditional on initial marital status suggest that married women might be more likely to adjust their labour supply in response to higher income of their spouses. This kind of within-household adjustment could be an important channel for the observed differences in labour market outcomes between men and women.

Figure 5 re-estimates equation (1) for the subsample of stable couples using both the log of individual income and the log of household income. Household income is measured as the sum of individual income and income of the spouse, but since it includes also individuals born outside the cut-off years of our main sample, the effects for individual men and women do not necessarily mirror each other.[19] For married men in Panel A, household income and individual income follow the same upward pattern over time. However, the estimated effect on household income is slightly smaller and is significant for a shorter time, such that from 7 years after the discovery of oil household income is not significantly different from zero at the 5% level. This suggests that, relative to low oil areas, wives of men in high oil regions earn less money than wives of men in low oil regions by 1980.

Panel B presents the same comparison for women. As seen previously, individual income

[19] On average, men in the sample tend to marry younger women and women in the sample tend to marry older men, so our household income is measured matching an individual to their spouse's income irrespective of whether the spouse is born between 1940–1949.

declines strongly for stably-married women in high oil regions and is significantly lower ten years after the discovery of oil relative to stably married women in low oil areas. However, a different picture emerges for household income. Women in high oil regions have higher levels of household income compared to women in low oil regions from 5 years after the discovery of oil. In the tenth year since the discovery, household income is 3.5% higher. Overall, while women's own income declines, the household as a whole seems better off after the discovery of oil.[20]

Given the significant fall in individual income and increase in household income, on average, our results reveal that women's economic standing in the household declines post-oil. To further illustrate this point, Figure 6 plots the annual fraction of own income in total household income in the years following the local shock, separately for men and women. While men's contribution to household income increases ten years after the oil discovery, it is not significantly different from that in 1969. In contrast, women's contribution to household income in high oil is 6% lower relative to the average fraction in 1969, an effect which is significant at the 5% level.

Overall, our results point to the importance of intra-household adjustments in labour supply following an economic boom as a mechanism behind the observed income declines among women. The reduction in own income is larger among married women, for whom household income increases. However, such an increase comes at the expense of within-household income equality. Given the evidence on the link between gender equality and declining domestic violence (Aizer, 2010; Anderberg et al., 2016) and the importance of the distribution of income within the household for marital satisfaction (Bertrand et al., 2015), it is unclear whether female and household welfare increases.

[20] Similar results are seen when relaxing the stable couple restriction, and calculating household income *irrespective* of whether the household remains together by 1980.

6.2 Robustness

In our baseline specification, we estimate the labour market impact of the oil boom by comparing outcomes of individuals who initially reside in areas which differentially become high/low oil. The lack of significant differences in the estimated coefficients prior to the discovery of oil in Figure 3 establishes the similarity of trends in income between high and low oil local labour markets pre-discovery. While the lack of diverging pre-trends lends support to the crucial parallel trends assumption, we address five potential identification concerns. First, we show that our baseline results are robust to *an alternative definitions of the oil treatment*, namely a direct spatial measure. Compared to the baseline definition, which relies on industry shares of oil in 1980, the spatial definition of oil abstracts from future developments and relies on the 'as good as random' nature of the location and size of the first oil discovery. Such a definition closely resembles that of Løken et al. (2012). Our results are not significantly different if we consider as treated the three closest areas to the initial Ekofisk discovery in 1969—Stavanger, Lister, and Haugesund in south-west Norway— and include in control areas other low-oil local labour markets (Figure S6) as well as all other local labour markets in Norway, including those defined as high- and mid-oil (Figure S7).

Second, we rule out *geographic mobility* as a driving force behind our results. Specifically, we examine how many individuals move, where they move to, and whether there are differences in geographic mobility between high and low oil regions (Section D.0.3) Importantly, results excluding movers are unchanged, suggesting that the above estimates are not driven by individuals who relocate (Figure S8.) In addition, geographic mobility is relatively low over the period under consideration, broadly similar between high- and low-oil, and the most common destination among movers in both high- and low-oil areas is Oslo, the low-oil capital.

Third, we show that our baseline results are not driven by the inclusion of *specific local* *labour markets* in a number of ways. To start with, we exclude Oslo and the Northern regions, most of which are defined as low-oil. Our results are similar, suggesting that particular local labour markets in the counterfactual group do not drive the observed changes in income in high-oil areas (Figures S9 and S10). We then consider a sample of only coastal local labour markets (both high- and low-oil), to show that differences in shipbuilding activities and coastal proximity do not drive the baseline results (Figure S11). Finally, we examine the importance of having a university or college in the local labour market. In particular, areas with higher education available may differ, for instance, with greater job opportunities available for women. Following a decree in 1965, local colleges expanded throughout the 1970s into areas which previously did not have opportunities for higher education (see for further details Carneiro et al., 2018). To examine this point, we define two separate samples: (1) local labour markets which had a university or college by 1980 and (2) local labour markets which did not have a university of college by 1980. Point estimates are virtually identical between the two subsamples (see Figure S12), and there is no evidence that the oil shock's impact made women in a local labour market that had a college or university lose less.

Our fourth identification concern is that *access to childcare* might differ between our treated and control areas. Substantial kindergarten expansions took place in Norway during the 70s and 80s, and these expansions happened faster in certain areas of the country (Havnes and Mogstad, 2011, 2015). To check if changes in kindergartens could be driving our results, we define the kindergarten expansion as in Havnes and Mogstad (2011, 2015). Figure S13 reveals that there is little correlation between the expansion of kindergarten and the fraction of employment in the oil industries in 1980, and differential exposure to kindergarten cannot explain the income changes due to the oil discovery.

Finally, we examine the potential role of religiosity in different parts of Norway as a possible driver of our results. Traditionally, the south-west of Norway is characterized by more conservative attitudes and higher religiosity. To see if areas of different religiosity, which may also be correlated with different conservative attitudes regarding women's role in the labour market, drive our results, we collect data on church membership of individuals in the 1960 Census. We construct a measure of religiosity at the local labour market level and combine it with our data on the fraction of employment in oil in 1980. We use two separate measures, to capture all types of religiosity: (i) membership in the state Church of Norway and (ii) membership in any other church, capturing all other religions apart from membership in the Protestant State Church.

Though there is geographic variation in religiosity, even the least religious areas have high levels of church membership. Importantly, there is no significant relationship between the fraction of people in a local labour market who are members of either the state or any other church and the fraction of employment in oil 20 years later. Religiosity is also not a key driver of our results: results are robust to excluding the top 25% most religious areas in 1960, suggesting the discovery of oil, rather than geographic differences in religion, are behind the observed income changes (see results in section D.0.4 in the Appendix).

6.3 Mechanisms

Our key finding is that the oil boom puts male and female workers onto diverging income trajectories. In particular, the finding is more pronounced among couples already married before the oil discovery. To understand what factors drive the decrease in income among women, we test the importance of some plausible mechanisms outlined in Section 2. First, changes in aggregate income depend on labour supply adjustments at the extensive margin—changes in labour force participation—and at the intensive margin—changes in full-time employment (Section 6.3.1).

Second, we explore sectoral effects in occupational dynamics to show that, while men benefit from the oil boom through increasing employment in high-paying occupations, women who remain in employment shift into other low-paying occupations (Section 6.3.2). Finally, we show that fertility and educational choices are largely unaffected by the oil discovery among women (Section 6.3.3). In the end, we provide some back-of-the-envelope calculations to identify which mechanisms might be more important.

6.3.1 Labour supply

We first consider changes in labour force participation, followed by changes in full-time employment. We supplement the annual Registry data with self-reported Census data from the two waves of 1970 and 1980 as described in Section 4.1.2. As before, we focus on young individuals born between 1940-1949 who are already of working age at the time of the oil discovery. We use a difference-in-differences framework to estimate the effect of the oil discovery on labour market outcomes by gender:

$$Y_{ibct} = \beta_0 + \beta_1 Oil_c + \beta_2 Post_t + \beta_3 Oil_c \times Post_t + \psi_c + \tau_b + \gamma_t + u_{ibct} \quad (2)$$

where Oil_c is the oil exposure developed in equation (1) and $Post_t$ is an indicator =1 in the year 1980.[21] The coefficient of interest is β_3, the interaction of the oil exposure with the post-discovery variable. Y_{ibct} corresponds to different labour market outcomes that allow us to examine different mechanisms. We explore both the extensive margin response, testing whether labour force participation changes, as well as the intensive margin response, testing for any change in full-time employment following the discovery of oil.[22]

[21] In the Census data, we take 1970 as the pre-period: since the discovery of the first giant oil field was at the very end of 1969, and production did not start until 1971, we can take 1970 census responses regarding labour outcomes as 'untreated' even in the high oil regions.

[22] In terms of income, taking Y_{ibct} as a measure of annual income, the Census data confirms what we found with

Labour market participation

First, we focus on the extensive margin. The main outcome variable is an indicator equal to one whenever an individual is not employed, either because he/she is outside of the labour force (inactive) or unemployed. Table 2 presents the results for male and female workers in odd and even columns respectively. Columns (1) and (2) present results for the sample of workers as defined in Section 4.2 who have strong attachment to the labour force prior to the oil discovery. Columns (3) and (4) present results for those with weak attachment to the labour force, those who were excluded from the sample in Section 4.2.[23]

The first column of Table 2 reveals that, following the oil discovery, men are less likely to leave the labour force in high oil regions compared to those in low oil regions. Relative to the average rate of non-employment in 1970, this represents a 25% decrease in labour force departure. Conversely, women in high oil regions are more likely to leave the labour force following the oil discovery relative to women in regions of low oil: non-employment increases by 29%. Similar patterns are observed for the sample of those who had weak attachment to the labour force from 1967–1969 (columns 3 and 4), and though the magnitude of the impact is smaller, it suggests that the local boom also draws some men into the labour force (12% compared to their pre-boom average) and discourages some women from entering the labour market (4%).

Full-time employment

Next, we examine whether there are labour supply adjustments at the intensive margin, conditioning on the sample of workers who are employed (for any amount of hours). Annual

the Registry annual data. Table S4 in Appendix E shows that men from the high oil regions gain around 12,000 NOK relative to men in the low oil areas by 1980, while women earn around 6,000 NOK less in 1980 compared to women in low oil areas.

[23]We define strong attachment to the labour force before the economic boom as having positive income in all years 1967–1969, as in our baseline results of section 5.2, while weak attachment is having zero income in at least one year from 1967–1969.

hours worked in the Census data are classified according to four bins: 1300 hours or more, 1000–1299, 500–999, and 100–499. We define a dummy for full-time employment equal to one for any amount of work entailing 1300 hours per year or more.

Table 3 shows there is little intensive margin response among men who are strongly attached to the labour force as full-time employment does not increase significantly following the discovery of oil. However, men who already work above our threshold of 1300 hours may have moved to work more hours, and we cannot rule out some degree of intensive margin adjustment within our broad category of full-time employment.

In contrast, women's probability of working full-time decreases significantly following the discovery of oil (by 6% relative to the mean in 1970). Such effects are even larger among women who are weakly attached to the labour force from 1967–1969: full-time employment declines by 11% among women in high oil. Thus, not only does the local boom lead to small decreases in full-time employment among women already working, but it also substantially discourages women in high oil from joining the labour force in full-time employment.[24] Note that, for our sample of young women in both low and high oil areas, there is a general tendency to reduce full-time employment by 1980, as shown by the negative and significant coefficient of $Post = 1$. Therefore, while all young women experience some degree of transition out of full time employment, possibly due to life-cycle events such as having children, the decline in full-time employment is even more pronounced for women in high-oil local labour markets.

6.3.2 Occupational dynamics

Changes in the labour supply already explain part of the income trends identified: fewer women in high oil regions work full-time and are more likely to leave the labour force, while men are

[24]Results in Appendix E.2 suggest that employment just less than full-time (1000–1299 hours annually) also decreases among women already working while employment with minimal hours worked (100–499 annually) increases.

less likely to leave the labour force. In addition, the discovery of oil may spur job creation in specific types of work. We examine which occupations grow between 1970–1980, which occupations contract, and whether changes in occupations among both men and women are concentrated among high or low paying occupations. Doing so sheds light on the changes in specific sectors that can cause the strong gender differences in income among high oil workers.

Figure 7 plots the changes in employment share between 1970 and 1980 by occupation, gender, and exposure to oil.[25] We only show the occupationswith at least a 0.5 percentage point change in employment shares over this period.[26] The occupations are ranked by the average income (from lowest to highest earning occupations). While there exists some overlap in the occupations that grow and contract, the set of occupations reported in Figure 7 tends to differ by gender. For instance, while women see changes in nursing and housekeeping, men see changes in engineering and manufacturing occupations for furnace, workshop, and cement workers.

Among men, Panel A shows that the share of workers employed in the lowest-paying jobs (such as agriculture and fishing) declines relatively more in high oil, while the share in higher-paying occupations (such as industrial and service jobs) increases in 1980 compared to 1970. Clearly, men reap the labour market benefits of the oil discovery: the increase in employment share is particularly stark in oil-related occupations in high oil relative to low oil markets, namely into engineering and technical jobs (the 21^{st} highest-paying job for men in 1980) or oil production jobs (the 9^{th} highest-paying job). This shift in the male labour force out of low-paying agricultural jobs towards high-paying oil-related occupations is one factor behind the boost in male income established earlier.

[25] Figure S18 plots the employment share of the same occupations in 1970. In high oil areas in 1970, the most common occupations for men are general farmers, workshop mechanics, farm-helpers, and cement finishers. For women, these are farm-helpers, shop assistants, other occupations in other clerical work, and charworkers. Relative to low oil areas, high oil areas have slightly higher levels of agriculture in 1970; Tables S2 and S3 reveal that pre-existing differences in the presence of agriculture do not explain any of the variation in high oil status.

[26] Overall, the figure shows 19 occupations for men, 24 occupations for women. This restriction is for visual convenience, as there are 288 occupations for men and 234 for women.

A different picture emerges for women. Panel B shows that between 1970 and 1980, women in high oil local labour markets move out of some of the low-paying, agricultural activities, such as farm-helpers, but the majority of women do not move into the oil sector or other high-paying jobs. Instead, women move into services and work as housekeepers (the 264^{th} highest-paying job among women in 1980), nurses (the 140^{th} highest-paying job), and secretaries (the 74^{th} highest-paying job). Overall, most women are not directly absorbed into the new, high-paying jobs created by the oil industry nor are occupational dynamics in spillover sectors conducive to income gains for women.

6.3.3 Education and fertility

Two other mechanisms could be relevant. First, given the fall in female income and the within-household reallocation of income shown in section 6.1, it is possible that families are making different fertility choices following the oil windfall. A plausible response to the positive income shock could be to have more children. Appendix G shows that this is indeed the case, and both women and man had higher chances of having a child and a higher number of children. However, the magnitude of the effect is small.[27] Combined with the within-household reallocation observed in Section 6.1, even the small observed changes in fertility are consistent with the fact that women shift to household production following the discovery of oil.

Secondly, we consider possible adjustments in education and training: the drop in income for young women could be driven by the choice to substitute work for further education, since the job opportunities for them are not so favourable immediately after the oil boom. Appendix I.3 shows that only men seem to invest more in schooling (although we find a small impact

[27]Figures S19 and S20 report our estimates. A decade after the oil discovery, the estimated difference in the number of children between high and low oil women is 0.248. Since the average number of children among women in low oil 10 years after the oil discovery is 1.72, our estimated difference represents a 14% increase, to 1.97 children, on average, per woman in high oil. The effect on the probability of having any child is even smaller, a 5% increase after 10 years.

of the oil shock on the years of education), while women do not significantly modify their educational attainments. These findings therefore rule out the possibility of women pulling away from the labour market for human capital accumulation as an explanation for their lower income trajectories following the oil boom.

6.3.4 Relative importance of different mechanisms

The previous sections examined the importance of different mechanisms that could explain the drop in female income following the oil boom. Having found evidence of changes in the female labour supply and in occupations, we would like to give some indication on which of these channels is predominant in explaining the loss in income. Due to the lack of data on hourly wage and on the exact number of hours worked, we can only approximately show the relative magnitude of these effects, looking at the transition between different labour statuses over the 1970s, namely (i) becoming unemployed, (ii) moving from full time to part-time and (iii) changing job (see Appendix H for details).

Overall, it seems that the main mechanism with respect to the labour supply and occupational dynamics identified in sections 6.3.1 and 6.3.2 is due to working women leaving the labour force. While we cannot map exactly how these transitions account for the income loss, since wages are also likely to change significantly in this time period, these results further point towards the importance of changes in the degree of employment among women.

6.4 Persistence: Lifetime Income of the Next Cohorts

In our baseline sample of workers aged 20–29 at the time of the discovery, women experience sizeable and persistent declines in their own income. Such losses are mitigated by increasing spousal income, at the expense of widening gender gaps within the household. Whether or not such a local economic shock causes continued decreases (increases) in the lifetime income of even younger women (men) has considerable implications for gender equality going forward. To further test the persistence of our results in the long run, we turn to a different sample: those who were too young to enter the labour force in 1969, at the start of the boom.

Younger cohorts are exposed to the shock in their formative years and are more capable of modifying key lifetime decisions compared to the baseline sample of workers aged 20–29. The cohorts born between 1952 and 1966 (aged four to 18 in 1970) grew up during the oil boom and had more flexibility in terms of job choice, education, and family formation decisions than their older peers. As a result, their labour market outcomes may differ compared to those already in the labour market during the economic boom.

Unlike our baseline estimations, we examine in this section the *lifetime* income of those who are exposed to the oil shock as children.[28] Following Haider and Solon (2006), we define lifetime income as the sum of an individual's income in the prime years of the average career path, namely between ages 30 to 45. Figure 8 reveals strong gender differences in the average lifetime income of nine birth cohorts residing in high/low oil labour markets prior to 1969. Younger men's lifetime income increases quickly in high oil, though this difference stabilizes across birth cohorts (panel A). Indeed, for men who experience the boom at even younger ages, the positive gap in lifetime income is remarkably stable across cohorts. In contrast, while older women at the time of oil discovery earn less than their low oil counterparts, the younger the women are at the time of the discovery, the smaller the negative gap in lifetime income (panel B). For the youngest cohort aged 4–6 at the time of oil discovery, the gap in lifetime income

[28]There are two advantages to using lifetime income over annual income when looking at long run impacts. First, the majority of our new sample is composed of children who are still in school and bound by compulsory schooling laws at the time of the shock. Thus, it is more relevant to consider their overall income rather than the income of a specific year. Second, lifetime income is free from any potential life-cycle biases which might result from differences across the treated and control regions due to the oil shock. Women and men in high oil labour markets may, for example, enter employment at younger ages and retire earlier, reflecting differential life-cycle decisions and not necessarily overall income gains.

between high and low oil areas vanishes.

These diverging patterns of lifetime income between men and women suggest that, while young women who are already working in booming regions have a permanent income loss after the oil shock, women in even younger cohorts have the time and opportunity to adapt to the changing economic conditions to avoid the adverse labour market outcomes that older cohorts experience. Formally, regression analysis of the lifetime income for different cohorts of children confirms this pattern (Appendix I). For males we find that, with the exception of the 13-15 age group, which has a marginally significant extra gain of 2% in lifetime income, the impact of the local boom for younger men in high oil is not statistically different from the impact for those aged 19-21. As seen in Figure 8, most boys experience a constant income gain independent of their birth cohort.

For females, in contrast, the estimated effect of being younger at the time of the boom is positive, large in magnitude, and statistically significant. Given that the lifetime income gap between high and low oil women aged 19–21 (the reference group) is -9.5%, the positive coefficients suggest that younger cohorts of women in high oil effectively close the lifetime income gap with women in low oil. Among women who were aged 13-18 in 1969, the lifetime income of those in high oil is only 5.2% less than those in low oil.[29] For the 4–12 age group, while the lifetime income gap remains significant, it is further reduced to 2.3–3.5%.

There can be several explanations for these results on the next generations. Other studies have found that early life conditions can have lifelong implications (see, for instance, Cunha and Heckman, 2007), and that an increase in family income can have especially large effects on children's outcomes for the lower end of the parental income distribution (Løken et al., 2012). However, it is challenging to isolate the importance of any one specific factor in explaining the long run recovery of income for young women who experience the economic boom in their

[29] This percentage results from -9.5+4.3%, the coefficient for this age bin from the second column of Table S8.

childhood. We present suggestive evidence that an increase in full-time employment across cohorts of women mirrors the observed catch up in lifetime income (Figure S24 in Appendix I, Panel B). We also show that migration (Figure S23) and decisions on human capital accumulation (Tables S9 and S10) seem to be less important factors.

7 Conclusion

In most developed economies, gender inequality in labour market outcomes persists despite years of policy initiatives. This article exploits the discovery of oil and gas in Norway to highlight how a large positive shock can trigger structural changes that hinder income equality and cause a temporary setback in women's labour market participation over a long time horizon. Using rich panel data on individuals and households, we find that men residing in regions most exposed to oil immediately benefit in terms of income and employment, while working women drop out of full-time employment and earn significantly less in the years following the boom. Furthermore, a decade after the oil discovery, a large share of men move into higher-paying occupations, while women do not.

Given that Scandinavian countries are consistently ranked among the most gender-equal in the world (WEF, 2020; UNDP, 2019), our results on the persistent income losses among women already in the labour force at the time of the oil discovery may appear surprising, particularly given the a-priori expectations of positive spillover effects of the oil boom on multiple sectors in the economy, but they are particularly important. These aggregate effects on young working-age women mask some interesting heterogeneity: women who are single at the time of the shock experience smaller declines in their individual labour income compared to married women. In terms of overall household income, however, married women experience moderate increases as their husbands' income rises. This result suggests a reallocation of the within-household labour supply, or women pulling back from the labour market, triggered by the higher male income.

A key novelty of our analysis lies in our ability to track the dynamics of income and labour supply decisions for the same individuals. Not only do we examine the effects and underlying mechanisms of the local shock on those already of working age at the time of the boom, but we also study the next generation of workers—those who were too young to be in the labour force during the first phases of the economic boom. We find that the youngest cohort of women do not suffer any lifetime income loss following the oil discovery, thus faring much better than women who are already in the workforce when the oil sector appears. The differences in income dynamics between women already in the labour force at the time of the discovery and younger ones who grow up during the boom suggest that flexibility in making career-related decisions in education, household formation, and labour supply decisions can have long-lasting effects on income. This result may indicate that there is room for timely policies to encourage labour supply adjustments that are more conducive to gender equality. It is important to note that many factors could be affecting the observed long term outcomes: for instance, the wealth brought by the local economic boom could catalyse socio-cultural changes that make it more acceptable for young girls to envision working full-time, even with working husbands, or for companies to demand female labour. A thorough exploration of these factors is beyond the scope of the current study.

Overall, our findings indicate that the oil boom is primarily a positive income shock for men, as women do not benefit economically from the arrival of oil and gas and its spillovers to other industries within the short to medium term. However, the losses for female workers are short-lived: in the long run, when new cohorts of women enter the labour market, the reduction in terms of income and labour market participation disappears.

Many important policy-relevant questions remain open for further analysis. For example, why do so few women switch into higher-paying jobs after the local boom unlike their male counterparts? What constraints do women face at the time, such as household duties, child-bearing and child care, and social taboos against women working in the oil sector? Lastly, what are the exact welfare consequences of this type of economic shock? We cannot evaluate in our study how changes in income translate into actual welfare experienced by these women. For example, it is possible that the intra-household reallocation of resources is welcome by both spouses, and that the reduction in annual income reflects an optimal labour supply choice that married women make willingly. It is also possible that the lower earning status is forced upon these women by the new economic context, and they do not have an option other than to work less, and/or for lower wages. This distinction remains to be validated in future research.

References

Acemoglu, D., D. Autor, and D. Lyle (2004). Women, war, and wages: The effect of female labor supply on the wage structure at midcentury. *Journal of Political Economy 112*(3), 497–551.

Aizer, A. (2010, September). The gender wage gap and domestic violence. *American Economic Review 100*(4), 1847–59.

Allcott, H. and D. Keniston (2018). Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. *The Review of Economic Studies 85*(2), 695–731.

Anderberg, D., H. Rainer, J. Wadsworth, and T. Wilson (2016). Unemployment and domestic violence: Theory and evidence. *The Economic Journal 126*(597), 1947–1979.

Aragón, F. M., J. P. Rud, and G. Toews (2018). Resource shocks, employment, and gender: Evidence from the collapse of the UK coal industry. *Labour Economics 52*, 54 – 67.

Arezki, R., V. A. Ramey, and L. Sheng (2017, 11). News shocks in open economies: Evidence from giant oil discoveries. *The Quarterly Journal of Economics 132*(1), 103–155.

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), 161–78.

Autor, D., A. Kostøl, M. Mogstad, and B. Setzler (2019, July). Disability benefits, consumption insurance, and household labor supply. *American Economic Review 109*(7), 2613–54.

Bartik, A. W., J. Currie, M. Greenstone, and C. R. Knittel (2019, October). The local economic and welfare consequences of hydraulic fracturing. *American Economic Journal: Applied Economics 11*(4), 105–55.

Basu, K. (2006). Gender and say: a model of household behaviour with endogenously determined balance of power. *The Economic Journal 116*(511), 558–580.

Benshaul-Tolonen, A. and S. Baum (2019). Structural transformation, extractive industries and gender equality. *OxCarre Working Paper N. 221*.

Bertrand, M., E. Kamenica, and J. Pan (2015, 01). Gender Identity and Relative Income within Households. *The Quarterly Journal of Economics 130*(2), 571–614.

Bhuller, M. S. (2009). Inndeling av norge i arbeidsmarkedsregioner. Technical Report Notater 2009/24, Statistisk sentralbyrå.

Black, D., T. McKinnish, and S. Sanders (2005). The economic impact of the coal boom and bust. *The Economic Journal 115*(503), 449–476.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016, February). Consumption inequality and family labor supply. *American Economic Review 106*(2), 387–435.

Brunstad, R. J. and J. M. Dyrstad (1997, 01). Booming sector and wage effects: An empirical analysis on Norwegian data. *Oxford Economic Papers 49*(1), 89–103.

Bütikofer, A., A. Dalla-Zuanna, and K. G. Salvanes (2018). Breaking the links: Natural resource booms and intergenerational mobility. Discussion paper no. 19/2018, NHH Dept. of Economics.

Carneiro, P., K. Liu, and K. G. Salvanes (2018, July). The Supply of Skill and Endogenous Technical Change: Evidence from a College Expansion Reform. IZA Discussion Papers 11661, Institute of Labor Economics (IZA).

Cascio, E. U. and A. Narayan (2015). Who needs a fracking education? The educational response to low-skill biased technological change. Working Paper 21359, National Bureau of Economic Research.

Cavalcanti, T., D. Mata, and F. Toscani (2019, March). Winning the oil lottery: the impact of natural resource extraction on growth. *Journal of Economic Growth 24*(1), 79–115.

Cools, S. and A. Kotsadam (2017). Resources and intimate partner violence in sub-saharan africa. *World Development 95*, 211 – 230.

Corden, W. M. and J. P. Neary (1982). Booming sector and de-industrialisation in a small open economy. *The Economic Journal 92*(368), 825–848.

Cunha, F. and J. Heckman (2007, May). The technology of skill formation. *American Economic Review 97*(2), 31–47.

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.

Donni, O. (2007). Collective female labour supply: theory and application. *The Economic Journal 117*(516), 94–119.

Ekeland, A. (2017). Employment in the petroleum industry and related industries 2016 - Sysselsatte i petroleumsnæringene og relaterte næringer 2016. *Statistics Norway Rapporter 2017/27*.

Fernández, R. (2013, February). Cultural change as learning: The evolution of female labor force participation over a century. *American Economic Review 103*(1), 472–500.

Fernández, R., A. Fogli, and C. Olivetti (2004). Mothers and sons: Preference formation and female labor force dynamics. *The Quarterly Journal of Economics 119*(4), 1249–1299.

Fleming, D. A. and T. G. Measham (2015). Local economic impacts of an unconventional energy boom: The coal seam gas industry in Australia. *Australian Journal of Agricultural and Resource Economics 59*(1), 78–94.

Goldin, C. (2006). The quiet revolution that transformed women's employment, education, and family. *American Economic Review 96*(2), 1–21.

Goldin, C. and C. Olivetti (2013). Shocking labor supply: A reassessment of the role of World War II on women's labor supply. *American Economic Review 103*(3), 257–262.

Goux, D., E. Maurin, and B. Petrongolo (2014, January). Worktime regulations and spousal labor supply. *American Economic Review 104*(1), 252–76.

Greenstone, M., R. Hornbeck, and E. Moretti (2010). Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. *Journal of Political Economy 118*(3), 536–598.

Haider, S. and G. Solon (2006, September). Life-cycle variation in the association between current and lifetime earnings. *American Economic Review 96*(4), 1308–1320.

Halla, M., J. Schmieder, and A. Weber (2019). Job displacement, family dynamics, and spousal labor supply. *American Economic Journal: Applied Economics forthcoming*.

Havnes, T. and M. Mogstad (2011). Money for nothing? universal child care and maternal employment. *Journal of Public Economics 95*(11), 1455–1465.

Havnes, T. and M. Mogstad (2015). Is universal child care leveling the playing field? *Journal of Public Economics 127*, 100 – 114. The Nordic Model.

Horn, M. K. (2011). Giant oil and gas fields of the world. *AAPG Datapages GIS Open Files*.

Jacobsen, G. D. and D. P. Parker (2014). The economic aftermath of resource booms: Evidence from boomtowns in the American West. *The Economic Journal 126*(593), 1092–1128.

James, A. and B. Smith (2017). There will be blood: Crime rates in shale-rich U.S. counties. *Journal of Environmental Economics and Management 84*(C), 125–152.

Jaworski, T. (2014). "You're in the army now:" the impact of World War II on women's education, work, and family. *The Journal of Economic History 74*(1), 169–195.

Kavanagh, M. (2013). Abandoned oil discoveries enjoy revival. *Financial Times*.

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.

Kotsadam, A. and A. Tolonen (2016). African mining, gender, and local employment. *World Development 83*, 325 – 339.

Larsen, E. R. (2006). Escaping the resource curse and the Dutch disease? When and why Norway caught up with and forged ahead of its neighbors. *The American Journal of Economics and Sociology 65*(3), 605–640.

Lei, Y.-H. and G. Michaels (2014). Do giant oilfield discoveries fuel internal armed conflicts? *Journal of Development Economics 110*, 139 – 157. Land and Property Rights.

Løken, K. V. (2010). Family income and children's education: Using the Norwegian oil boom as a natural experiment. *Labour Economics 17*(1), 118 – 129.

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

Lundberg, S. (1985). The added worker effect. *Journal of Labor Economics 3*(1), 11–37.

Marchand, J. (2012). Local labor market impacts of energy boom-bust-boom in Western Canada. *Journal of Urban Economics 71*(1), 165–174.

Maurer, S. E. and A. V. Potlogea (2020). Male-biased demand shocks and women's labour force participation: Evidence from large oil field discoveries. *Economica*.

Michaels, G. (2011). The long term consequences of resource-based specialisation. *The Economic Journal 121*(551), 31–57.

Ngai, L. R. and B. Petrongolo (2017, October). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics 9*(4), 1–44.

Norwegian Petroleum (2020, Updated: 12.05.2020). Norway's petroleum history. *Ministry of Petroleum and Energy and Norwegian Petroleum Directorate*. https://www.norskpetroleum.no/en/framework/norways-petroleum-history/.

Olivetti, C. and B. Petrongolo (2016). The evolution of gender gaps in industrialized countries. *Annual Review of Economics 8*(1), 405–434.

Parkes, K. R. (2012). Shift schedules on north sea oil/gas installations: A systematic review of their impact on performance, safety and health. *Safety Science 50*(7), 1636 – 1651.

Sandbu, M. (2009). The Iraqi who saved Norway from oil. *Financial Times* (August 28).

Stephens, Jr., M. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics 20*(3), 504–537.

Toews, G. and P.-L. Vezina (2017). Resource discoveries and FDI bonanzas; an illustration from Mozambique. OxCarre Working Papers 199, Oxford Centre for the Analysis of Resource Rich Economies, University of Oxford.

UNDP (2019). Beyond income, beyond averages, beyond today: Inequalities in human development in the 21st century. Technical report, Humand Development Report.

Van Den Bark, E. and O. D. Thomas (1981, 11). Ekofisk: First of the Giant Oil Fields in Western Europe. *AAPG Bulletin 65*(11), 2341–2363.

WEF (2020). Global gender gap report 2020. Technical Report ISBN-13: 978-2-940631-03-2, World Economic Forum, Switzerland.

Wilson, R. (2020, apr). Moving to economic opportunity: The migration response to the fracking boom. *Journal of Human Resources*, 0817–8989R2.

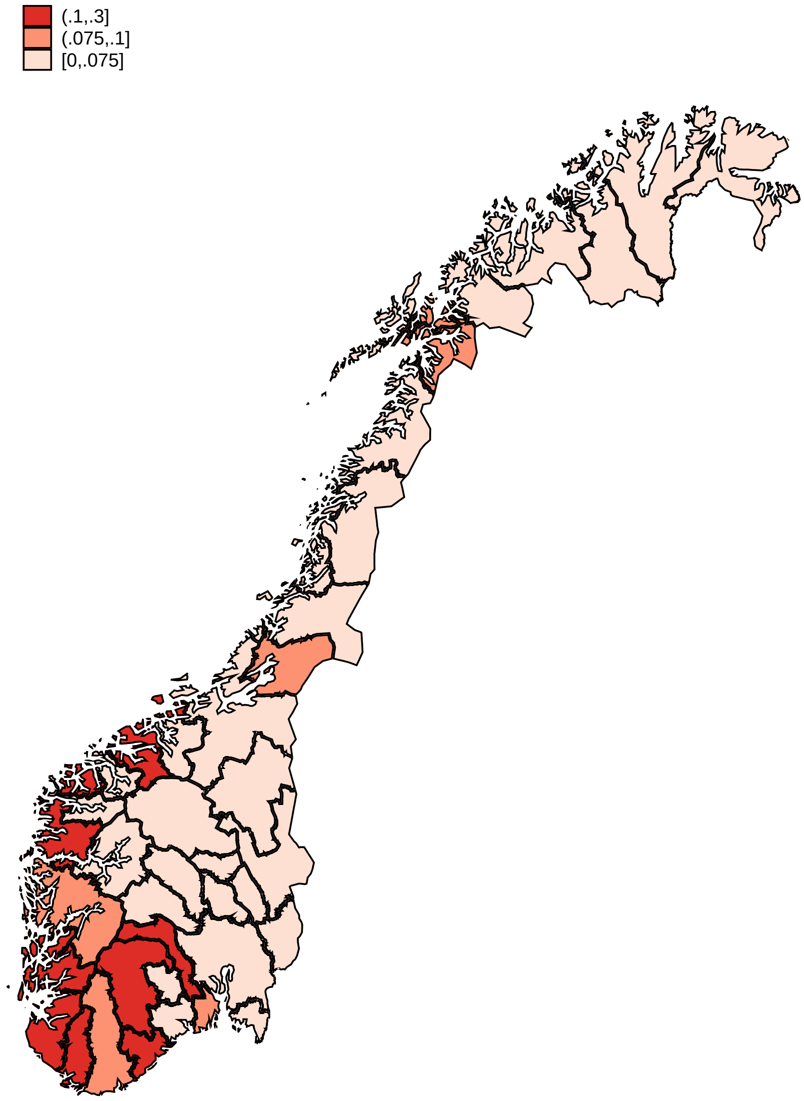

Figure 1: Geographic Distribution of the Influence of the Oil Discovery. Exposure to oil is defined as the share of employment in an oil industry or secondary industry supporting oil production in each local labour market. High oil (dark red): fraction of employment in oil larger than 0.1. Medium oil (orange): fraction of employment in oil between 0.075 and 0.1. Low oil (light pink): fraction of employment in oil less than 0.075.

Panel A. Men

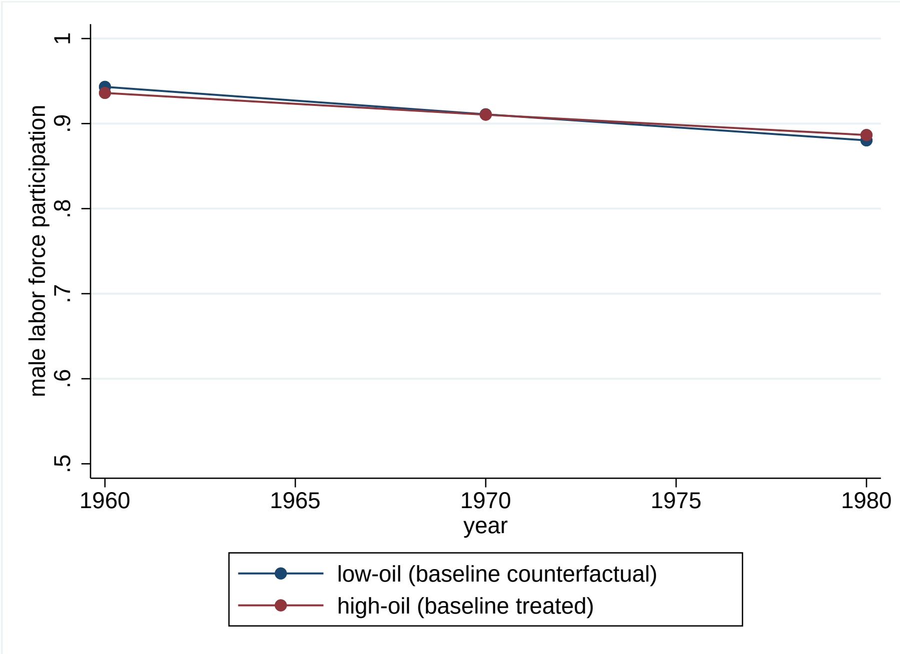

Panel B. Women

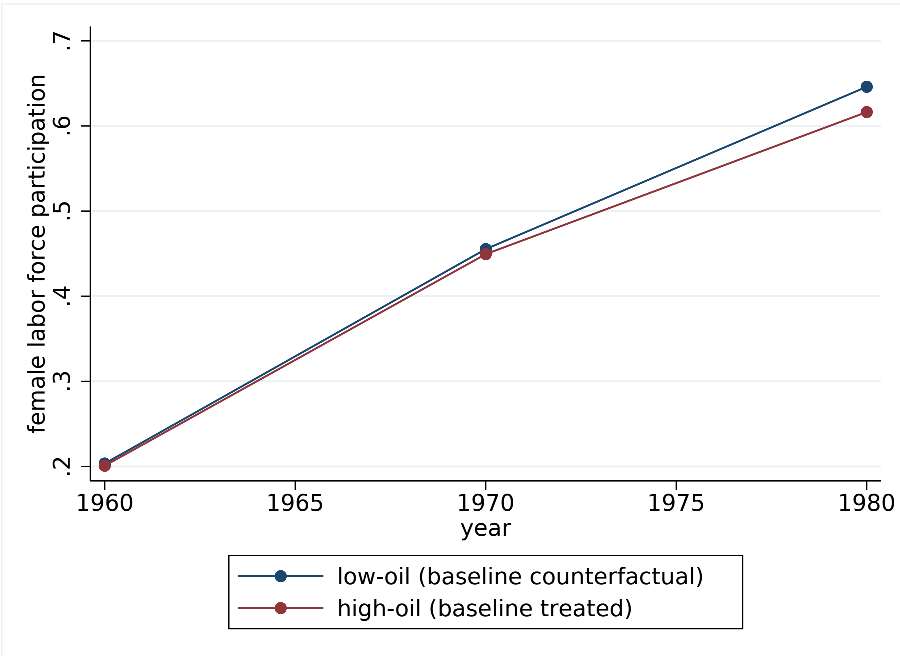

Figure 2: Evolution of labour force participation by oil status, ages 18–64. The figure plots the male and female labour force participation rate in high-oil and low-oil local labour markets in 1960, 1970 and 1980. Note that while values of y-axis differ between panels, the scale of the y-axis is the same, such that lines are directly comparable.

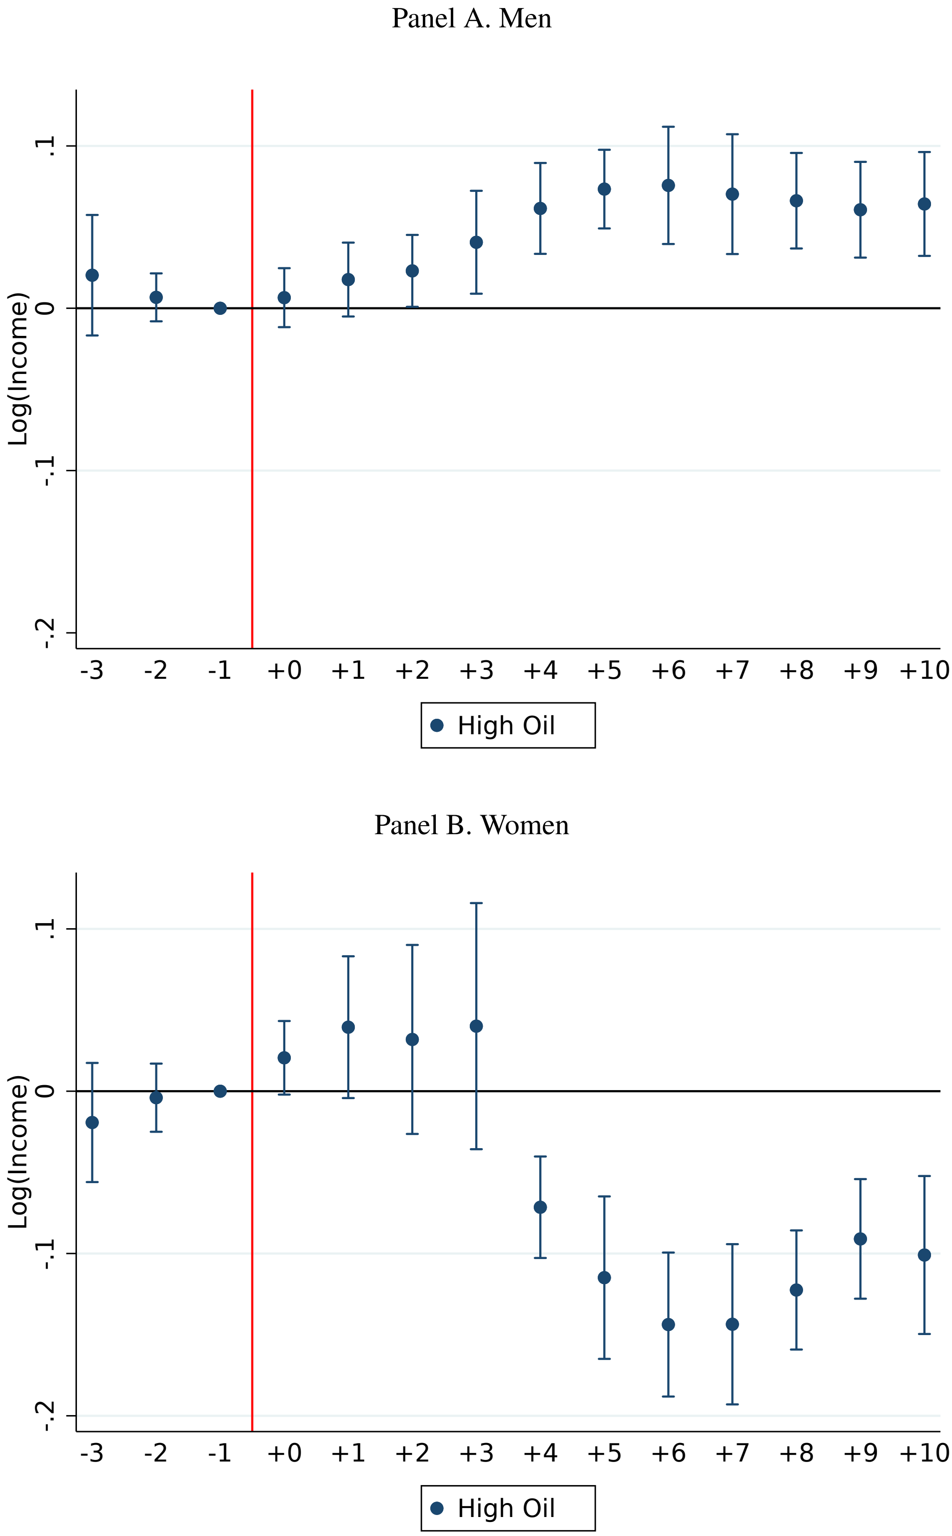

Figure 3: Estimated Effects of the Oil Discovery on Log(Income) over Time by Gender. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on annual log income by gender, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated in a low oil local labour market in 1965 in equation (1) from separate regressions by gender. The red vertical line corresponds to the discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

Panel A. Men

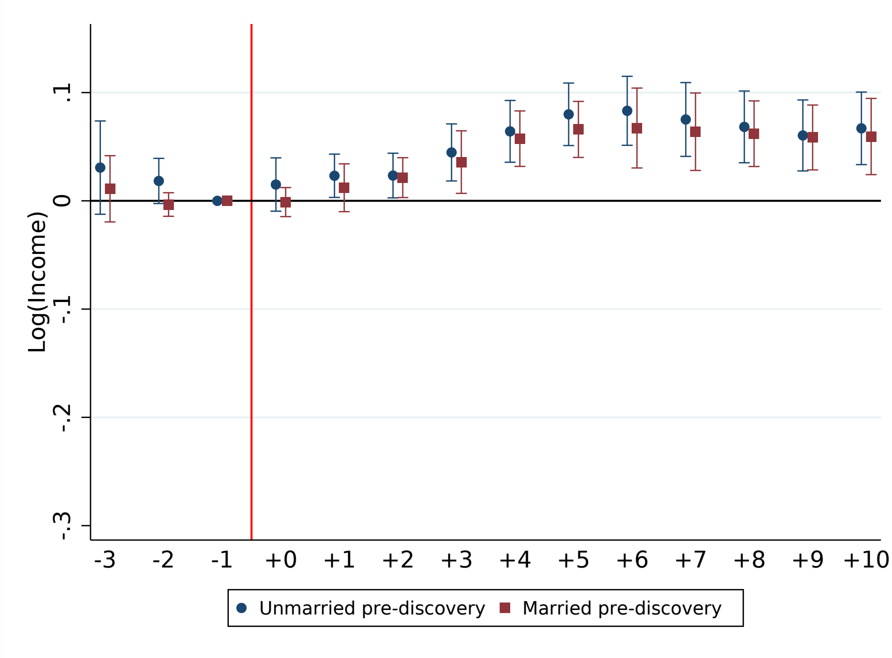

Panel B. Women

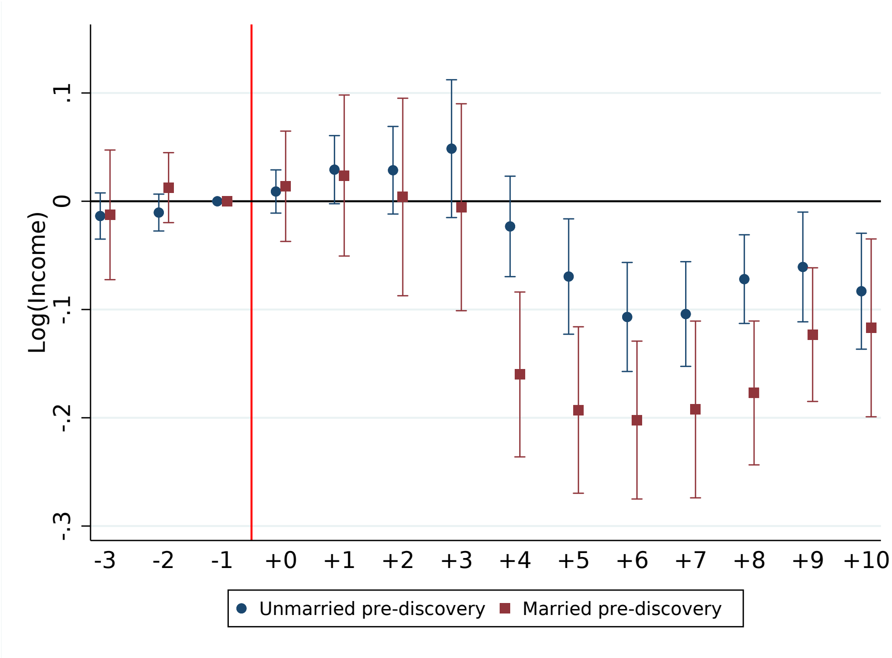

Figure 4: Estimated Effect of Oil Discovery on Log(Income) over Time by Gender and Marital Status. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on annual log income by gender and initial marital status, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated ones in low oil local labour markets in 1965 in equation (1) from separate regressions by gender and marital status. Initial marital status is defined prior to the oil shock in 1969. The red vertical line corresponds to the discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

Panel A. Men

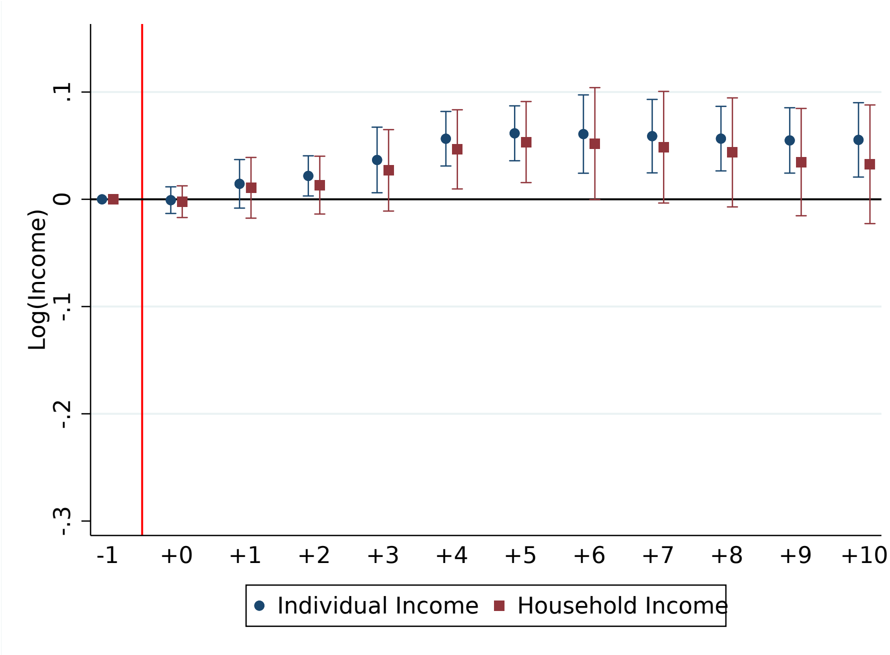

Panel B. Women

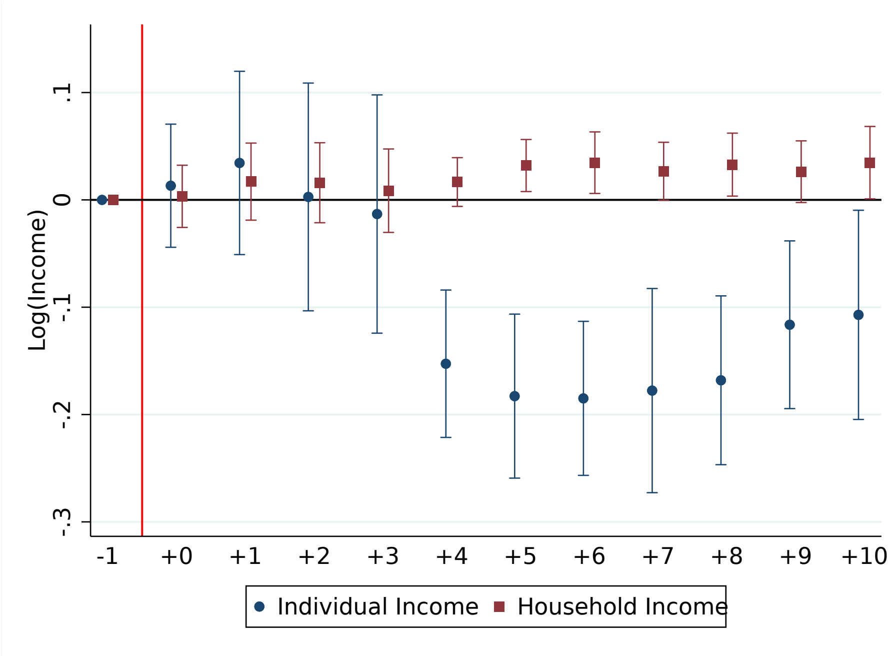

Figure 5: Estimated Effect of Oil Discovery on Individual and Household's Log(Income) over Time by Gender. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on annual log income for single individuals and for their household by gender, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated ones in low oil local labour markets in 1965 in equation (1) from separate regressions by gender. Individual and household regressions are estimated using the same sample of stable couples, defined as those who are married in 1969 (or earlier) through to 1980. The trends for men and women do not need to sum to the household trend, because the household income measure includes also marriages to individuals outside the sample birth cohorts. Time horizon spans -1 to $+10$ as marital status is defined in 1969, thus some couples were unmarried prior to -1. Stable couples represent 38% and 35% of the sample in Figure 3 among men and women respectively. The red vertical line corresponds to the discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

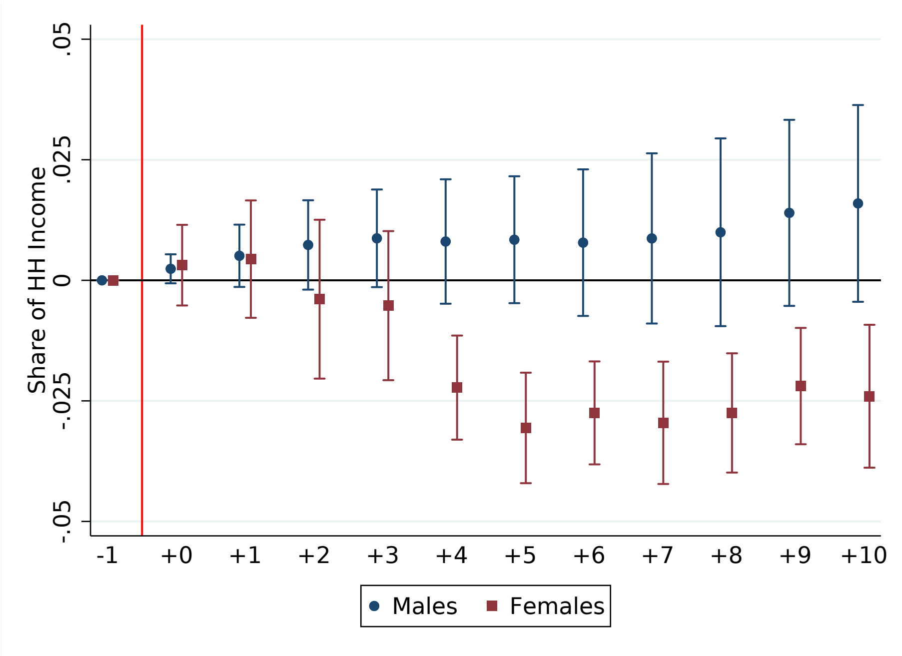

Figure 6: Estimated Effect of Oil Discovery on Individual Share of Household Log(Income) Over Time by Gender. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on the male and female share of their household's annual log income, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated ones in low oil local labour markets in 1965 in equation (1) from separate regressions by gender. Regressions are estimated using the same sample of stable couples defined as those who are married in 1969 (or earlier) through to 1980. Time horizon spans -1 to $+10$ as marital status is defined in 1969, thus some couples are unmarried prior to -1. Stable couples represent 38% and 35% of the sample in Figure 3 among men and women respectively. The red vertical line corresponds to the discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero. Average share of household income in 1969 is 0.86 and 0.41 for men and women respectively.

Panel A. Men

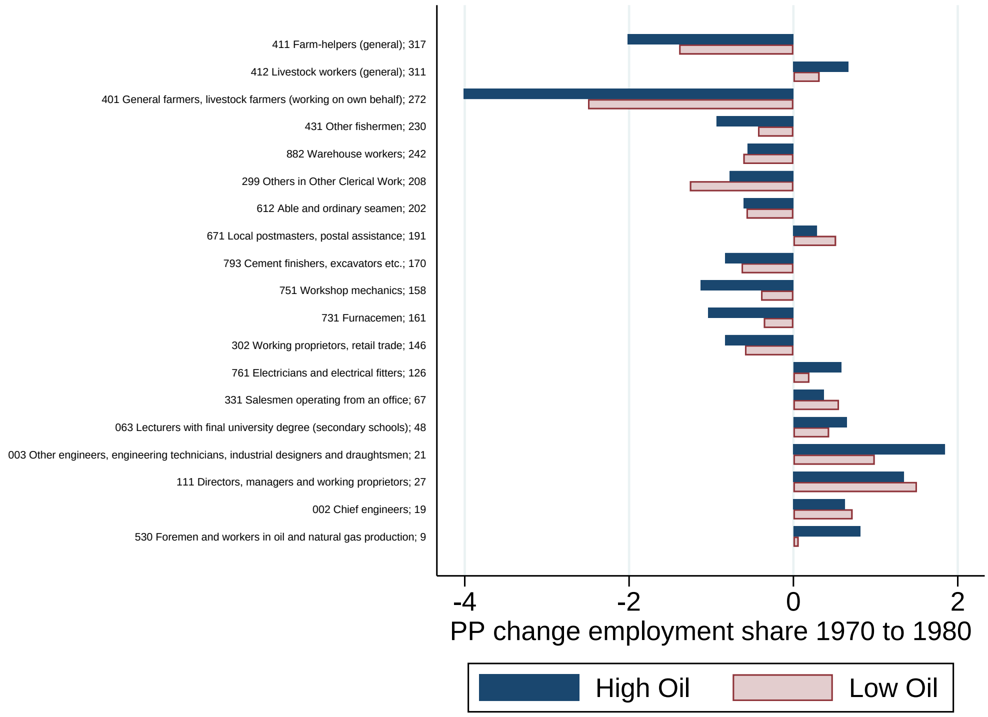

Panel B. Women

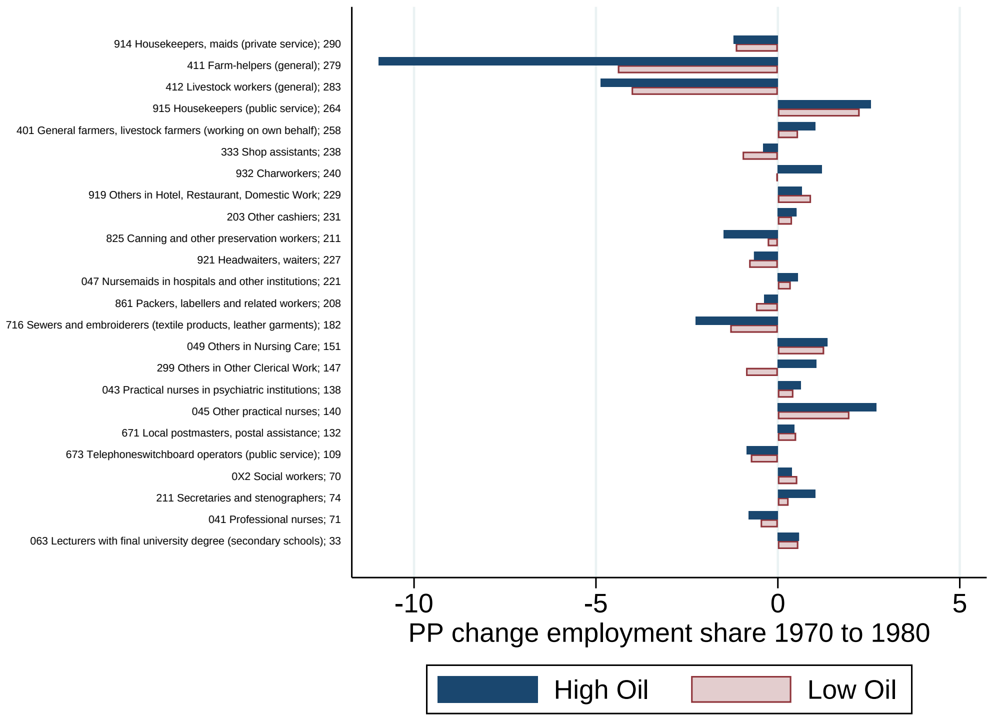

Figure 7: Change in Employment Shares between 1970 and 1980. The figure plots the percentage point changes in employment shares in high versus low oil local labour markets between 1970 and 1980, among occupations whose employment share increased/decreased by at least 0.5 percentage points. The number at the end of each occupation denotes the national rank of the occupation in terms of average income earned in 1980. Occupations classified according to the International Classification of Occupation (ISCO) system. Sample of all workers in high/low oil employed in an occupation in a given year.

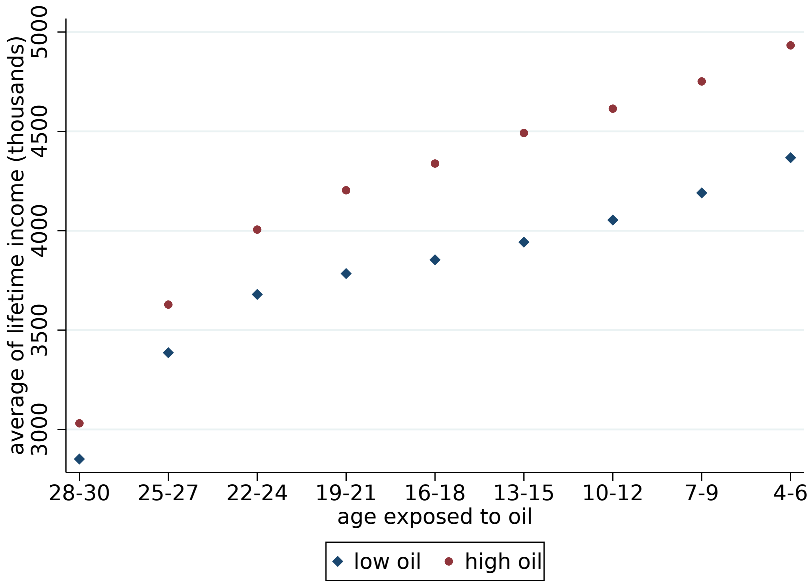

Panel A: Men

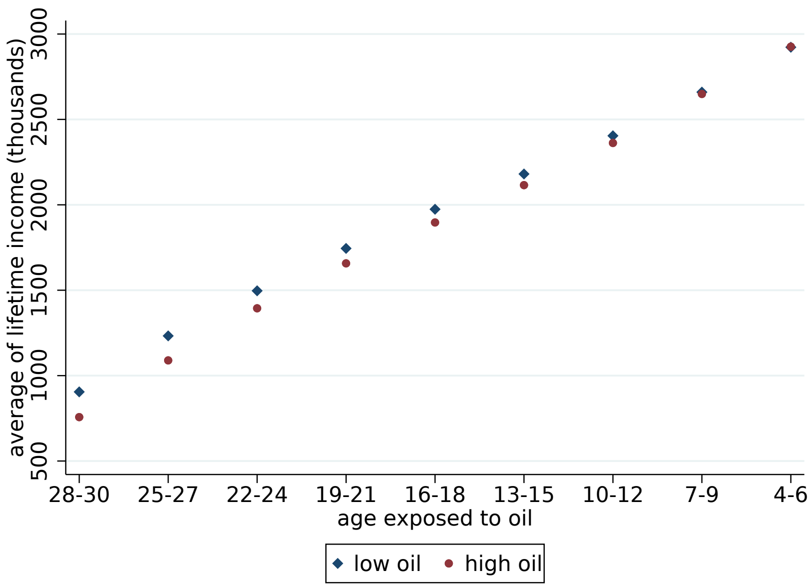

Panel B: Women

Figure 8: Lifetime Income Gap by Age Cohort. The figure plots the average lifetime income (measured between ages 30–45) of groups of people who had different ages at the time of oil discovery in 1970. Each cohort corresponds to a three-years bin, such that the 4-6 cohort was born in 1966-1964, and aged 4-6 in 1970; the 7-9 cohort was born in 1963-1960, and so on. The sample includes all individuals born between 1939 -1966. High/low oil defined as residing in high/low oil local labour markets in 1965 (from birth for those born after 1965). Income measured in NOK, inflation adjusted to year 1998.

Table 1: Summary Statistics of Young Workers (Birth Cohorts 1940-1949) in 1969

	(1) Male	(2) Female	(3) Both Genders
Pension income (in 1998 NOK)	151,006.0	105,745.2	136,430.1
	(66139.5)	(47373.5)	(64309.2)
Age	24.66	23.80	24.39
	(2.730)	(2.619)	(2.725)
Living in High Oil LLM in 1965	0.167	0.147	0.160
	(0.373)	(0.354)	(0.367)
Living in Medium Oil LLM in 1965	0.188	0.186	0.187
	(0.391)	(0.389)	(0.390)
Living in Low Oil LLM in 1965	0.645	0.668	0.652
	(0.478)	(0.471)	(0.476)
Married	0.611	0.603	0.608
	(0.488)	(0.489)	(0.488)
Observations	182,570	86,724	269,294

Notes: Table reports means and standard deviations of demographic characteristics for men and women with positive income in the three years before the oil discovery as described in Section 4.2.

Table 2: Census Results - Not Employed

Sample:	Strongly Attached 67–69		Weakly Attached 67–69	
	(1) Not Emp. Male	(2) Not Emp. Female	(3) Not Emp. Male	(4) Not Emp. Female
Oil LLM=1	0.004***	-0.028***	0.001	-0.015**
	(0.001)	(0.007)	(0.007)	(0.006)
Post=1	0.014***	0.112***	-0.139***	-0.367***
	(0.001)	(0.005)	(0.009)	(0.003)
Oil LLM=1 × Post=1	-0.006***	0.050***	-0.025*	0.029**
	(0.002)	(0.013)	(0.013)	(0.013)
Initial LLM FE	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Observations	310938	149304	134306	270732
R^2	0.003	0.034	0.047	0.142
Mean Dep. Variable 1970	0.024	0.170	0.209	0.659

Notes: Table reports the difference-in-differences estimates of equation (2) separately by gender using a sample that includes birth cohorts between 1940-1949 (aged 21-30 in 1970). Dependent variable is a dummy variable indicating non-employment, defined as either outside the labour force or unemployed. Treatment is high exposure to oil based on 1965 residence. 'Strongly attached' to LF defined as having positive income in all years 1967–1969, while 'weakly attached' is having zero income in at least one year from 1967–1969. Standard errors clustered at the initial local labour market level are reported in parentheses. ***$\rho < 0.01$, **$\rho < 0.05$, *$\rho < 0.10$.

Table 3: Census Results - Full-Time Employment

Sample:	Strongly Attached 67–69		Weakly Attached 67–69	
	(1) Full Time Male	(2) Full Time Female	(3) Full Time Male	(4) Full Time Female
Oil LLM=1	-0.029***	0.011**	-0.024**	0.041**
	(0.005)	(0.005)	(0.010)	(0.017)
Post=1	0.063***	-0.275***	0.304***	-0.031**
	(0.006)	(0.009)	(0.018)	(0.013)
Oil LLM=1 × Post=1	0.006	-0.045***	0.001	-0.050*
	(0.010)	(0.013)	(0.019)	(0.025)
Initial LLM FE	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Observations	293879	112497	113303	136652
R^2	0.029	0.085	0.136	0.014
Mean Dep. Variable 1970	0.852	0.733	0.591	0.422

Notes: Table reports the difference-in-differences estimates of equation (2) separately by gender using a sample that includes birth cohorts between 1940-1949 (aged 21-30 in 1970). Dependent variable is a dummy variable indicating full-time employment, defined as being employed 1300 hours or more annually. Sample is restricted to those with any employment in 1970 and 1980. Treatment is high exposure to oil based on 1965 residence. 'Strongly attached' to LF defined as having positive income in all years 1967–1969, while 'weakly attached' is having zero income in at least one year from 1967–1969. Standard errors clustered at the initial local labour market level are reported in parentheses. ***$\rho < 0.01$, **$\rho < 0.05$, *$\rho < 0.10$.

Online Appendix

A Giant Oil Fields in Norway

Figure S1: Giant offshore oilfields in Norway. The map shows the geographic distribution of giant oilfields discovered in Norway in the 1970s and 1980s, all situated offshore. For each cluster of discoveries, we report the name and year of only one field. Ekofisk, located south-west of Stavanger in the North Sea, was the first of these discoveries and the one used for our empirical analysis.

B Breakdown of Oil Industry in 1980

Table S1: Employment Shares of Detailed Industries Comprising Oil, Young Workers (Birth Cohorts 1940-1949)

	(1) Male High Oil	(2) Male Low Oil	(3) Female High Oil	(4) Female Low Oil
emp share manufacturing of machinery (including oil rigs)	0.0511	0.0255	0.0116	0.00479
	(0.220)	(0.158)	(0.107)	(0.0690)
emp share manufacturing of transport equipment	0.0795	0.0305	0.0113	0.00320
	(0.270)	(0.172)	(0.106)	(0.0565)
emp share non-building construction (including oil well drilling)	0.0682	0.0491	0.00605	0.00571
	(0.252)	(0.216)	(0.0776)	(0.0754)
emp share petroleum/natural gas production, petroleum refinery and manufacturing	0.0107	0.00302	0.00148	0.000726
	(0.103)	(0.0549)	(0.0385)	(0.0269)
Observations	30710	118931	8756	44096

Notes: Table reports employment shares for detailed industries comprising the definition of oil described in Section 5.1 in 1980 for men and women with positive income in the three years before the oil discovery as described in Section 4.2 who also have an industry classified in 1980.

C Income in Levels and Probability of Positive Income

Panel A. Men

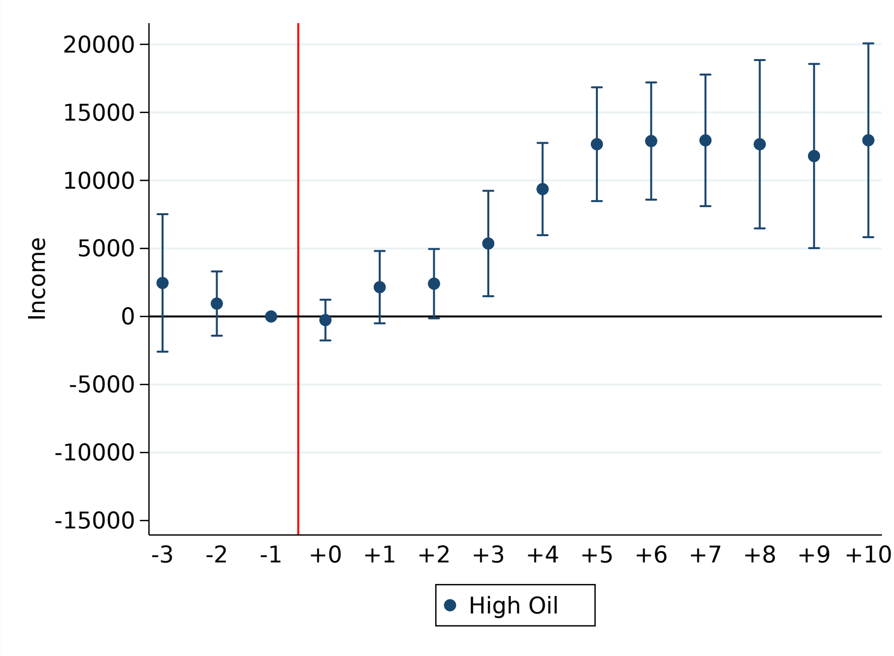

Panel B. Women

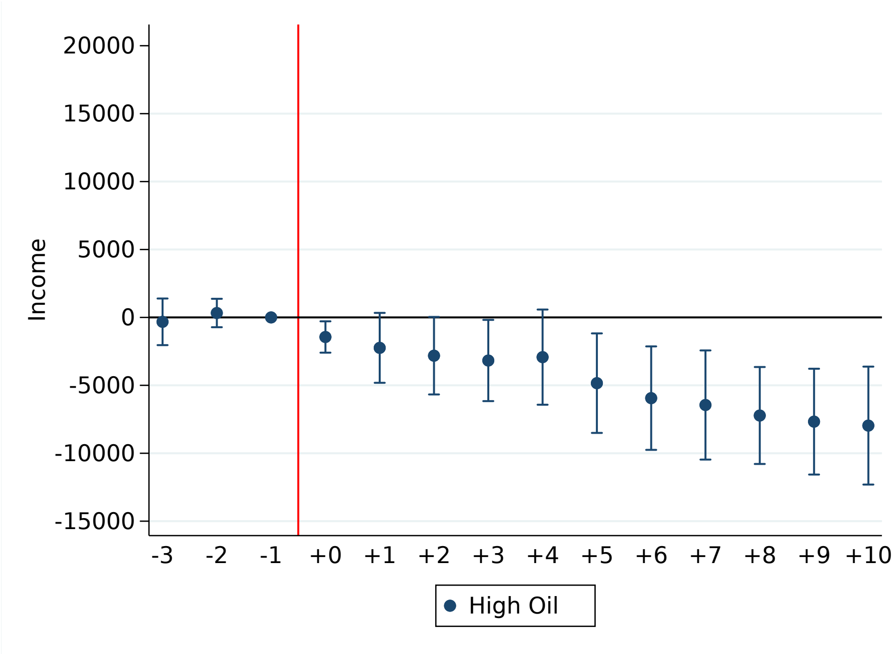

Figure S2: Estimated Effects of the Oil Discovery on Income over Time by Gender. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on annual income by gender, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated ones in a low oil local labour markets in 1965 in equation (1). The red vertical line corresponds to discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

Panel A. Men

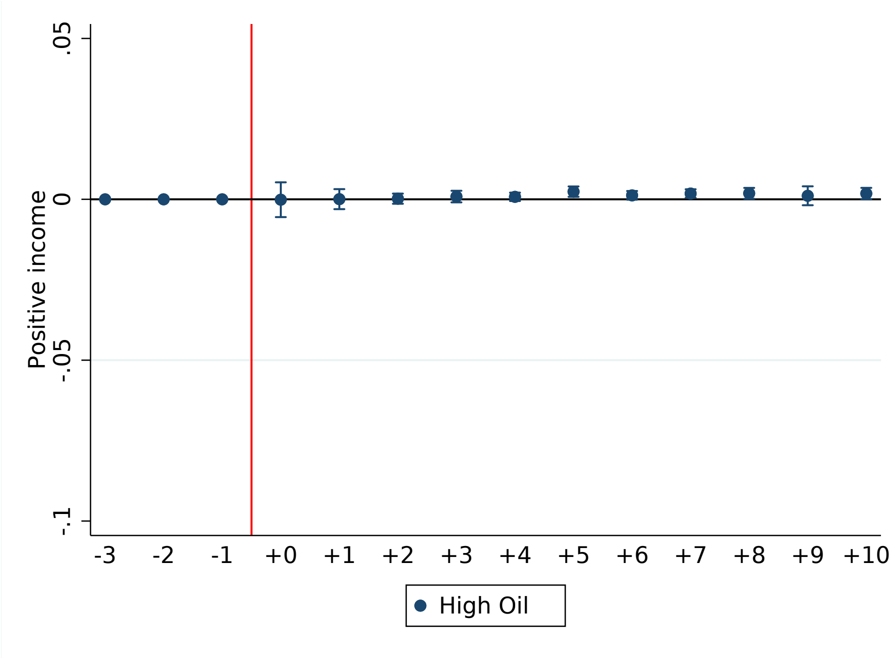

Panel B. Women

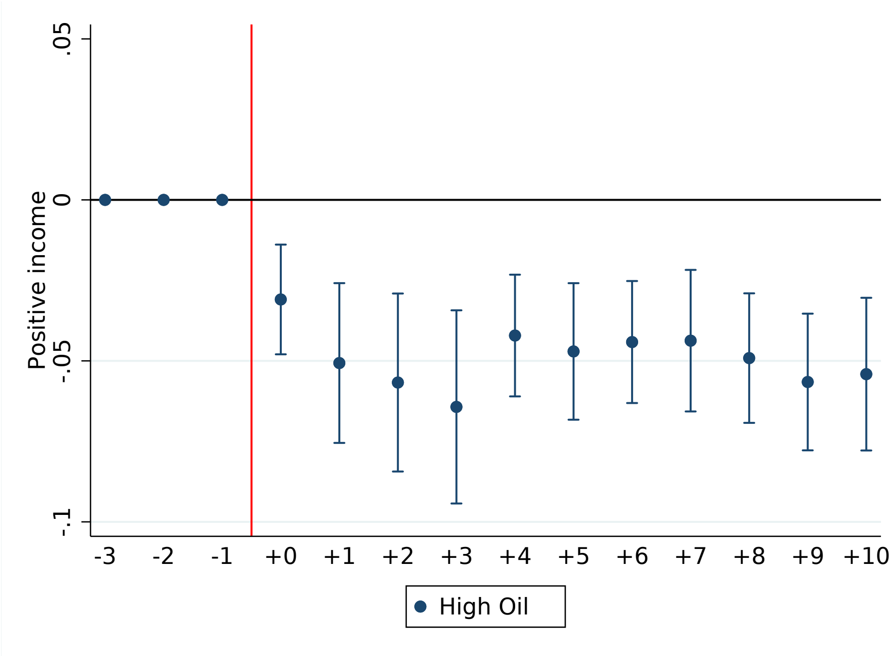

Figure S3: Estimated Effects of the Oil Discovery on the Probability of Having Positive Income over Time by Gender. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on the probability of having positive income by gender, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated ones in a low oil local labour markets in 1965 in equation (1). The red vertical line corresponds to discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

Panel A. Men

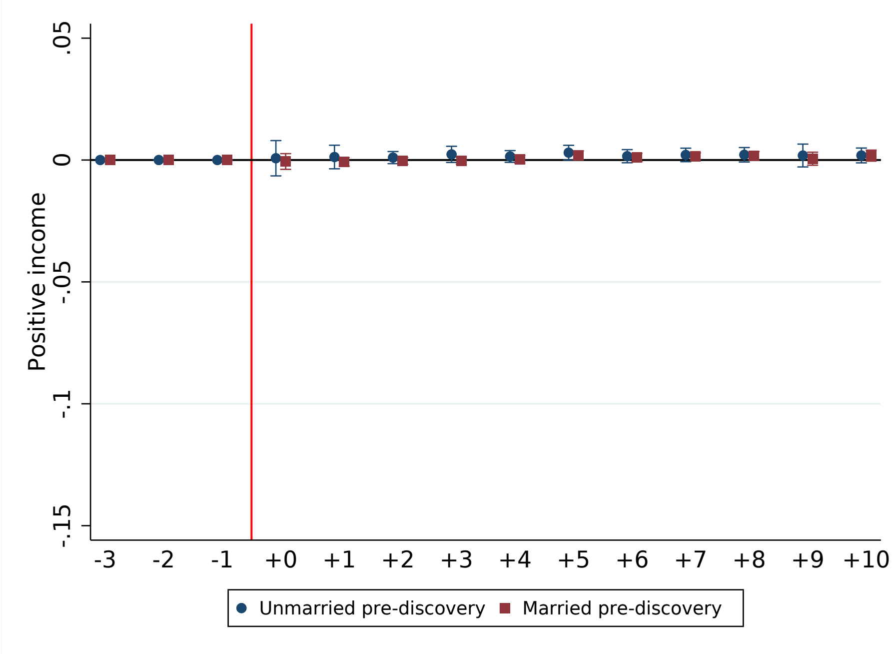

Panel B. Women

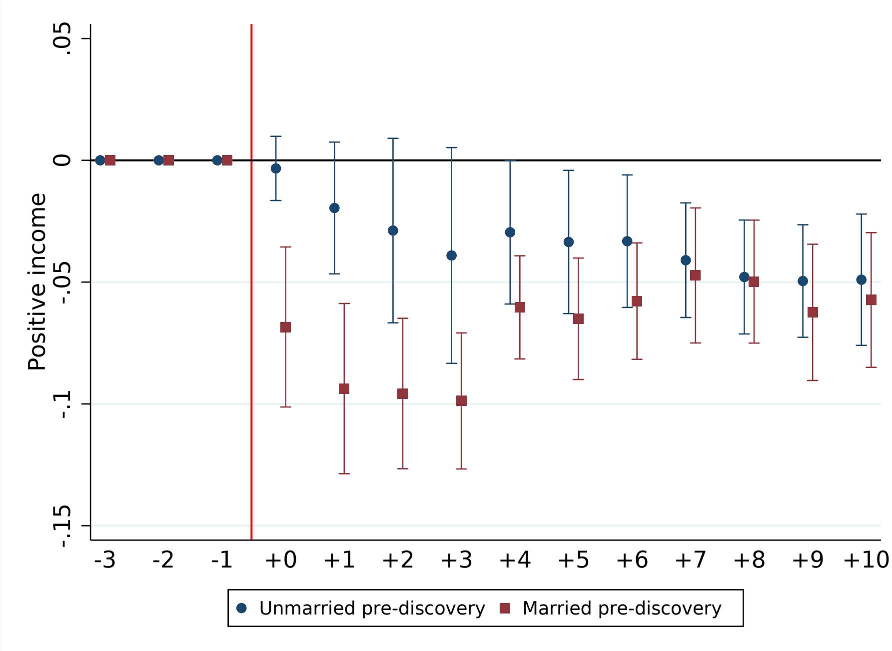

Figure S4: Estimated Effect of Oil Discovery on the Probability of Having Positive Income Over Time by Gender and Marital Status. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on the probability of having positive income by gender and initial marital status, comparing treated individuals who reside in a high oil local labour market in 1965 to untreated ones in low oil local labour markets in 1965 in equation (1) from separate regressions by gender and marital status. Initial marital status is defined prior to the oil shock in 1969. The red vertical line corresponds to the discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

D Robustness

D.0.1 Can Pre-Existing Characteristics Explain High Employment Share in Oil in 1980?

Our definition of high oil regions is based on the characteristics of labour markets after the economic boom, namely the share of employment in oil in 1980. While the difference-in-differences approach does not require the variation in treatment intensity be unrelated to local labour market characteristics, it is useful to examine the determinants of the expansion of oil employment across local labour markets up to the 1980s. Table S2 reports results of separate regressions of high oil status in 1980 on 2-digit industry shares in 1960, examining whether the most common industries predict high oil status in 1960. Of the industries examined, the employment share in the manufacture of metal products, machineries and equipment in 1960 is positively related to high oil in 1980 and is the most significant predictor of high oil status. Included in this 2-digit industry is shipbuilding, which also benefits from the discovery of offshore oil. Despite this, industry and the other five most common industries that accounted for more than half of total employment before the boom can only predict around 12% of the likelihood of becoming a high oil local labour market (see the R^2 of Table S2).

Even with the inclusion of other controls, such as employment in agriculture and the public sector, population density, educational characteristics and a control for coastal municipalities in Table S3, the ability of these local labour market-specific characteristics to predict which areas become high oil after the boom does not increase significantly. Male employment in metal products, machineries and equipment manufacturing remains one of the few significant determinants. Figure S5 details the robustness of the results to accounting for differences in this industry, illustrating that such industries do not drive our results as annual income for men and women still evolves as in our baseline specification. Overall, these results suggest that oil development by 1980 is largely unrelated to pre-existing area factors.

Table S2: Predicting High Oil Status in 1980 using 1960 Census

	(1) High Oil 1980
emp. share agricultural	-0.0962
	(-0.19)
R^2	0.000
emp. share transportation and storage	3.162*
	(1.69)
R^2	0.079
emp. share construction	-2.661
	(-0.85)
R^2	0.025
emp. share social and community services	0.815
	(0.19)
R^2	0.000
emp. share retail trade	-6.225
	(-0.95)
R^2	0.020
emp. share manufacture of metal products, machinery, and equipment	3.624***
	(2.75)
R^2	0.119
LLMs	46

Notes: Each panel corresponds to separate regression of 6 most common 2-digit industries on the probability of a local labour market to become high oil. All industries account for 53% of total employment in 1960. The last industry does have a lot of predictive power on its own, since the manufacture of equipment includes shipbuilding, one of the activities that most strongly predicts the future development of a large oil industry. Table regresses binary variable equal to 1 indicating high oil measured in 1980 (as defined in Section 5.1) on local labour market-level industry shares in 1960 before the discovery of oil. Estimating equation: $highoil_c = \beta_0 + \beta_1 empshare_c + u_c$.

Table S3: Predicting High Oil Status in 1980 using 1960 Census

	Female			Male		
	(1) High Oil	(2) High Oil	(3) High Oil	(4) High Oil	(5) High Oil	(6) High Oil
employment share, manufacturing	0.979	0.502	0.233	0.023	-0.015	-0.996
	(1.870)	(1.920)	(2.182)	(1.262)	(1.651)	(1.405)
employment share, public sector	-1.183	-0.901	-1.072	-3.929	-2.700	-3.620
	(2.239)	(2.617)	(2.677)	(3.148)	(3.568)	(3.094)
employment share, agriculture and fishery	0.866	-0.696	-0.717	-0.759	-0.519	-0.653
	(1.965)	(3.964)	(3.988)	(1.173)	(1.383)	(1.399)
fraction of population with any occupation		-2.768	-2.721		-15.836	-9.404
		(2.042)	(2.040)		(14.890)	(14.447)
fraction of total population female/male		-1.592	-1.851		8.233	7.340
		(6.441)	(6.573)		(10.464)	(9.527)
total population density		42.641	43.150		55.742	42.980
		(41.016)	(42.301)		(34.141)	(37.529)
LLM non-costal indicator		-0.060	-0.091		0.067	-0.026
		(0.188)	(0.176)		(0.227)	(0.206)
fraction of population low educated		1.426	1.538		0.867	1.581
		(1.464)	(1.467)		(1.666)	(1.539)
employment share, manufacture of metal products, machinery, and equipment			1.191			3.633**
			(3.675)			(1.790)
LLMs	46	46	46	46	46	46
R^2	0.088	0.201	0.204	0.055	0.228	0.302
F-statistic, joint significance	1.533	1.675	1.565	1.178	1.724	2.071
p-value, joint significance	0.219	0.131	0.163	0.329	0.119	0.059

Notes: Table regresses binary variable equal to 1 indicating high oil measured in 1980 (as defined in Section 5.1) on local labour market-level factors in 1960 before the discovery of oil. Factors in 1960 are gender specific (with the exception of population density and non-coastal indicator). As such the fractions take, for example, the number of women (men) employed in an industry divided by the total number of women (men) in the local labour market. Standard errors are clustered at the local labour market level.

Panel A. Men

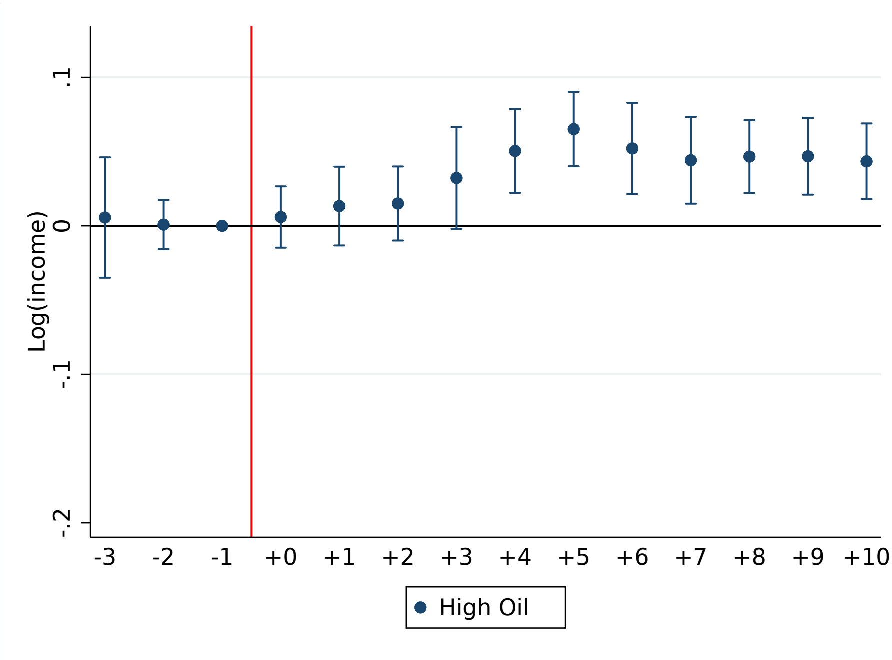

Panel B. Women

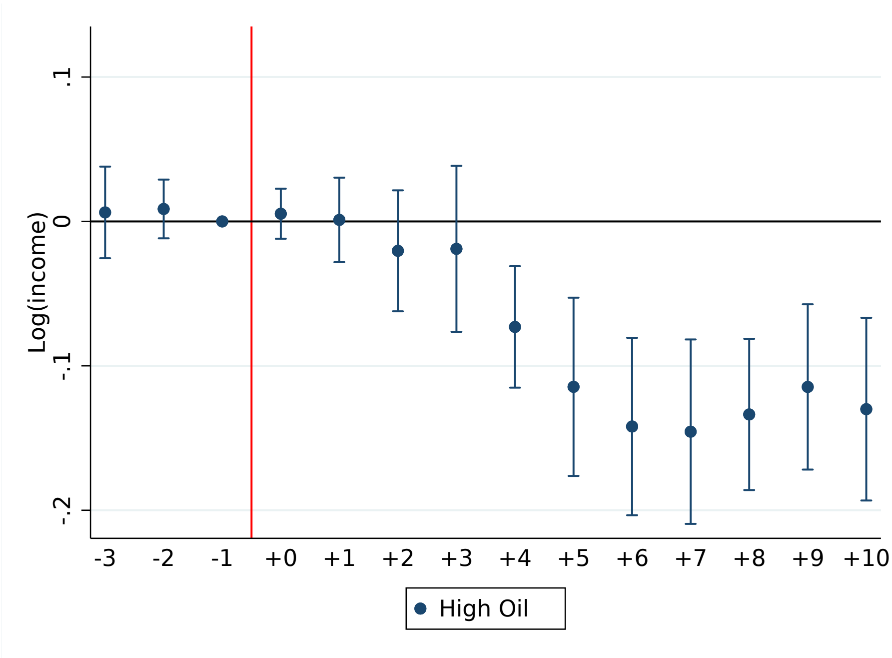

Figure S5: Estimated Effect of the Oil Discovery on Log(Income) over Time by Gender Excluding some Secondary Oil Industries. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on annual log income for individuals as in baseline specification of Figure 3 excluding local labour markets with high share of employment in manufacturing of metal, machinery, and equipment in 1960, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated ones in low oil local labour markets in 1965 in equation (1). Estimation sample—high and low oil areas, excluding any local labour markets with more than 10% of employment in MFG of metal, machinery, and equipment in 1960 both in high and low oil. The red vertical line corresponds to the discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

D.0.2 Alternative Definitions of Oil Local Labour Markets

Figures S6 and S7 redefine treated local labour markets to exploit the direct spatial measure of treatment by oil. Compared to the baseline definition, which relies on employment shares in oil in 1980, exploiting the direct spatial definition of oil abstracts from post-oil discovery development. Rather, the spatial measurement of oil relies on the as good as random nature of the location of the first oil discovery. This definition compares the three closest areas to the initial Ekofisk discovery in 1969—the local labour markets of Stavanger, Lister, and Haugesund—to other local labour markets in Norway.

Figure S6 compares these three local labour markets to other low oil local labour markets, a definition similar to Løken et al., 2012, while Figure S7 compares these three local labour markets to *all* local labour markets, including those defined as high- and mid-oil in the baseline definition. Reassuringly, results for both men and women are unchanged compared to Figure 3, though they are of slightly smaller magnitude in Figure S7 when the counterfactual group has been expanded substantially.

Panel A. Men

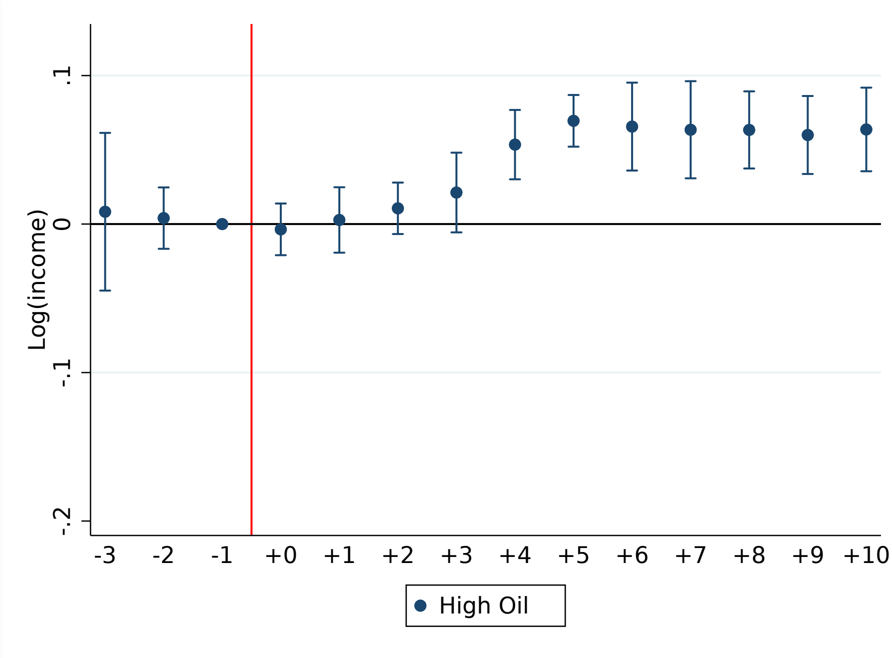

Panel B. Women

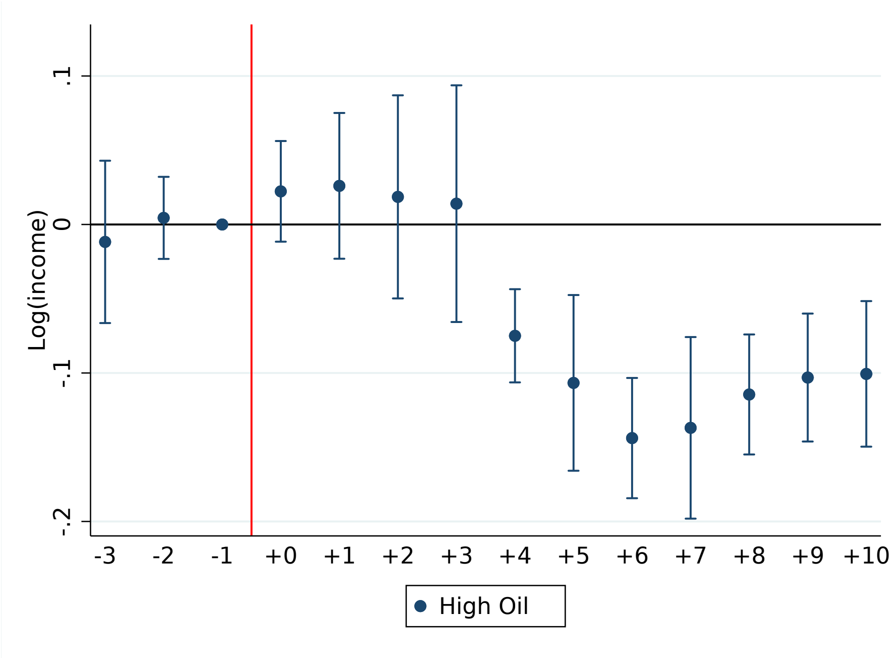

Figure S6: Estimated Effect of the Oil Discovery on Log(Income) over Time by Gender with Spatial Identification. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on annual log income for individuals as in baseline specification of Figure 3, but comparing the three high-oil local labour markets closest to initial discovery (Stavanger, Lister, and Haugesund, all of which become high-oil) as treated to all low-oil ones as controls. This design exploits the spatial component of the discovery rather than subsequent developments and closely mirrors Løken et al. (2012). The red vertical line corresponds to the discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

Panel A. Men

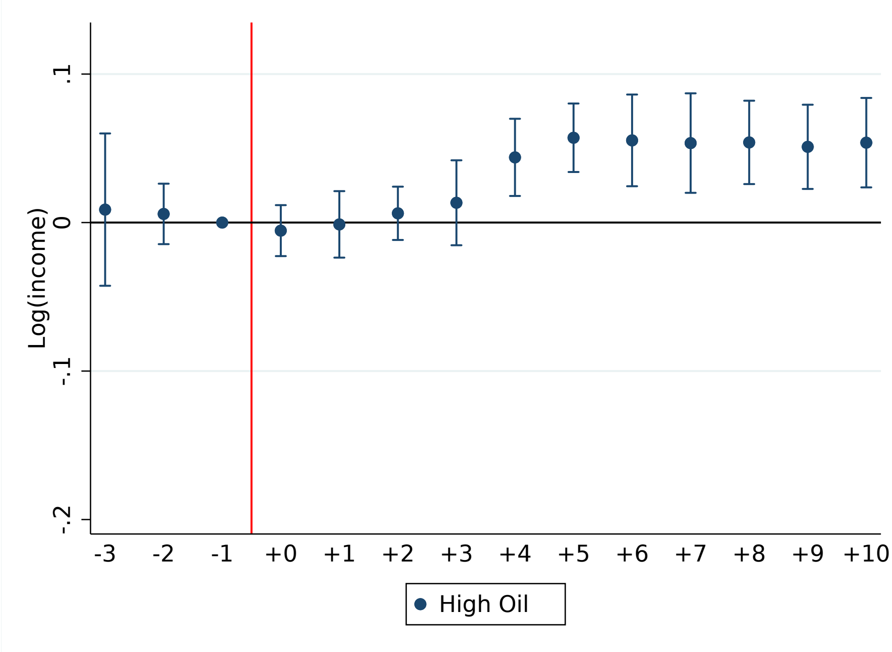

Panel B. Women

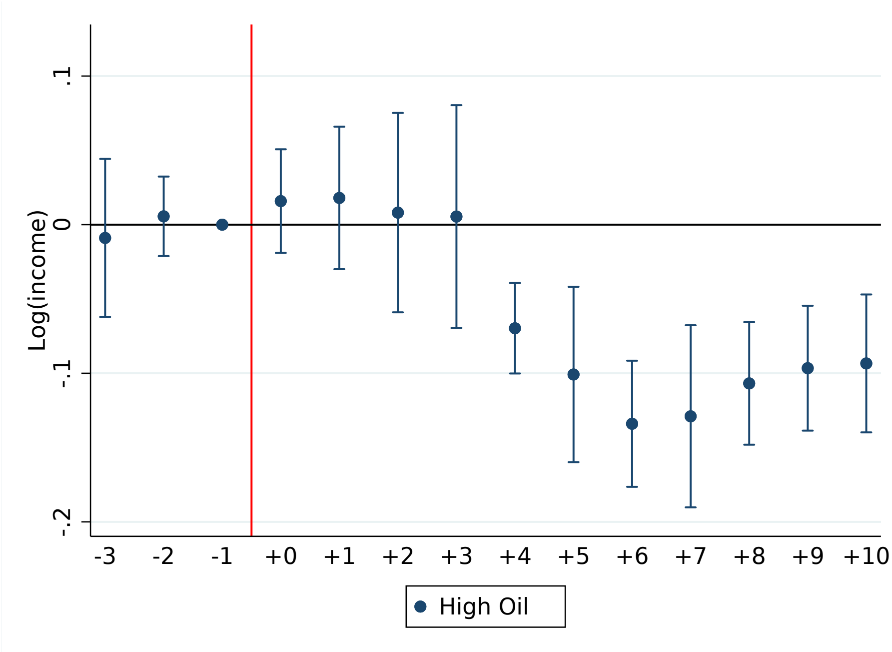

Figure S7: Estimated Effect of the Oil Discovery on Log(Income) over Time by Gender with Spatial Identification - Different Control Group. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on annual log income for individuals as in baseline specification of Figure 3, but comparing the three high-oil local labour markets closest to initial discovery (Stavanger, Lister, and Haugesund, all of which become high-oil) as treated to all low-oil and mid-oil ones as controls. This design exploits the spatial component of the discovery rather than subsequent developments and closely mirrors Løken et al. (2012). The red vertical line corresponds to the discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

D.0.3 Geographic Mobility

Our baseline results could be underestimating the changes in the local workforce resulting from migration into and out of the oil regions. Constructing the treatment and control groups using an individual's local labour market of residence *prior* to the oil discovery partly limits concerns about endogenous relocation of workers following the change in economic circumstances. However, if a large fraction of the population moved in response to the oil shock, then the average income in the treatment and control groups would be an inappropriate comparison since many individuals would no longer be living in the initial location. Indeed, the literature often identifies some geographic mobility following resource discoveries. For instance, Wilson (2020) shows significant migration into oil regions following the fracking boom.

Differential rates of geographic mobility do not appear to be problematic for the analysis. In our baseline sample, the fraction of workers who move is similar between those in high and low oil local labour markets. Among men, 21% of those in high oil and 23% of those in low oil reside in a new local labour market in 1980 compared to 1965. Among women, 35% of those in high oil and 30% of those in low oil regions reside in a new local labour market in 1980.

While geographic mobility is relatively low over the period, and roughly similar between high and low oil local labour markets, mobility patterns among those who do move may differ between high and low oil region movers. Among men who migrate and were initially residing in a high oil local labour market, 25% relocate to another high oil local labour market and 51% relocate to a low oil local labour markets. Among those who migrate and who were initially residing in a low oil local labour market, 70% relocate to another low oil local labour market and 13% relocate to a high oil local labour market. As just 23% of low oil area men move to a new local labour market by 1980, only 3% of those who initially resided in a low oil local labour market reside in a high oil local labour market in 1980. The most frequent destination among both high and low oil movers is Oslo, the financial centre of the country and a low-oil area. 23% of those who move from high oil local labour markets and 25% of those who move from low oil local labour markets ultimately reside in Oslo in 1980. This suggests that mobility into high oil regions among men is not a major concern.

Geographic mobility is higher among women relative to men, and slightly higher among women from high oil local labour markets. As with men, the most frequent destination among both high and low oil movers is Oslo, where 25% of high oil movers and 27% of low oil movers ultimately settle. As with men, only 4% of those who initially resided in a low oil local labour market reside in a high oil local labour market in 1980, suggesting that mobility into high oil regions is not a major concern.

Additionally, Figure S8 examines how the results of Figure 3 change when excluding movers. If differential mobility out of or into high oil local labour markets strongly impacts the results of Figure 3, then excluding these movers would provide a different picture. Reassuringly, Figure S8 reveals that results are consistent when excluding movers.

Panel A. Men

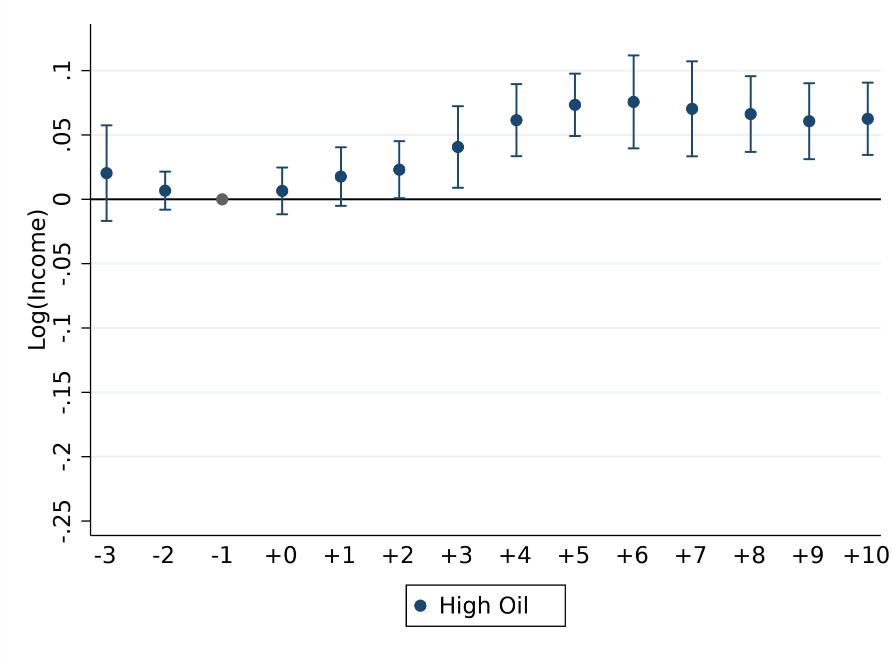

Panel B. Women

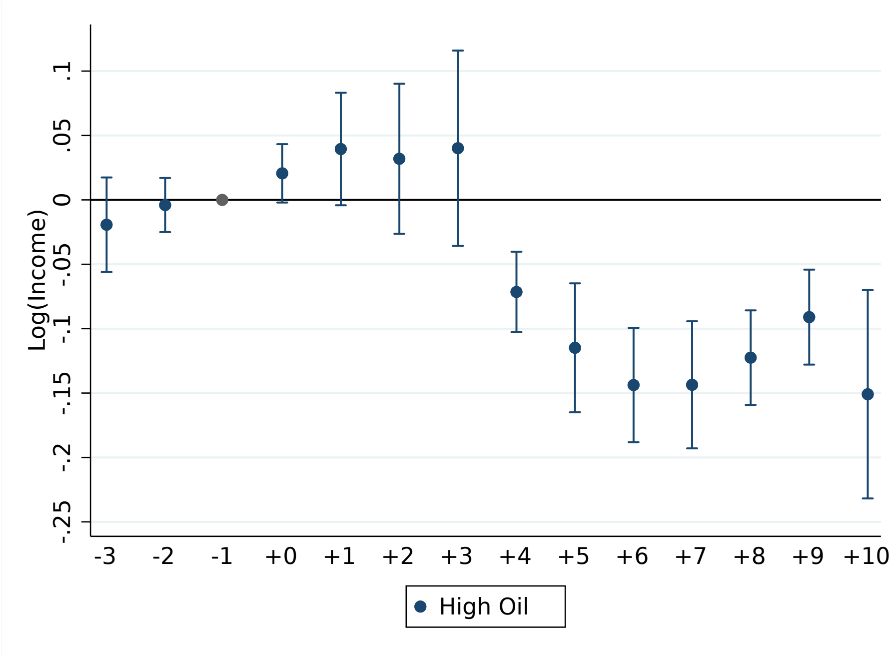

Figure S8: Estimated Effects of the Oil Discovery on Log(Income) over Time by Gender Excluding Movers. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on log of annual income by gender, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated in a low oil local labour market in 1965 in equation (1). The estimation sample excludes those who reside in a different local labour market from where they resided in in 1965. Vertical line corresponds to discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

D.0.4 Do Specific Local Labour Markets Drive the Results?

We implement three robustness checks to examine how influential specific local labour markets are in the baseline results. The baseline results are unchanged if: (i) we exclude the largest local labour market, Oslo, the capital, from the control group which had 4.7% of employment in oil in 1980 (Figure S9); (ii) we exclude the Northern local labour markets from the sample, most of which are low-oil (Figure S10); (iii) we consider only coastal municipalities, excluding landlocked areas that were farther from the offshore oil discovery (Figure S11); and (iv) we exclude labour markets that experienced substantial college expansions (Figure S12). The stability of these results suggests that particular local labour markets do not drive the observed changes in income in high- versus low-oil areas. Considering a sample of comparable areas in terms of coastal access, together with the results of Figure S5, suggests that differences in shipbuilding and coastal proximity do not drive the baseline results, nor do the access to educational facilities.

Panel A. Men

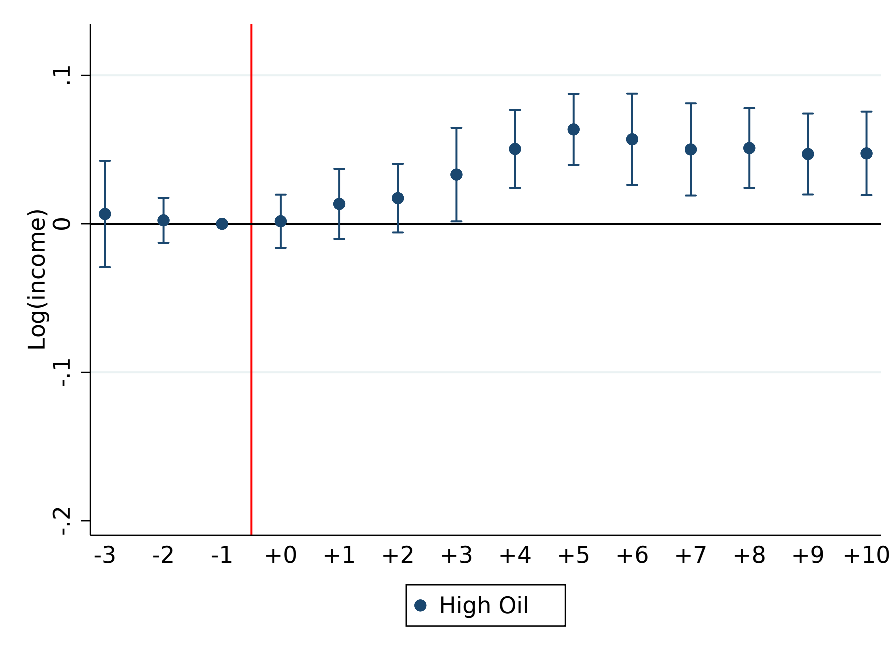

Panel B. Women

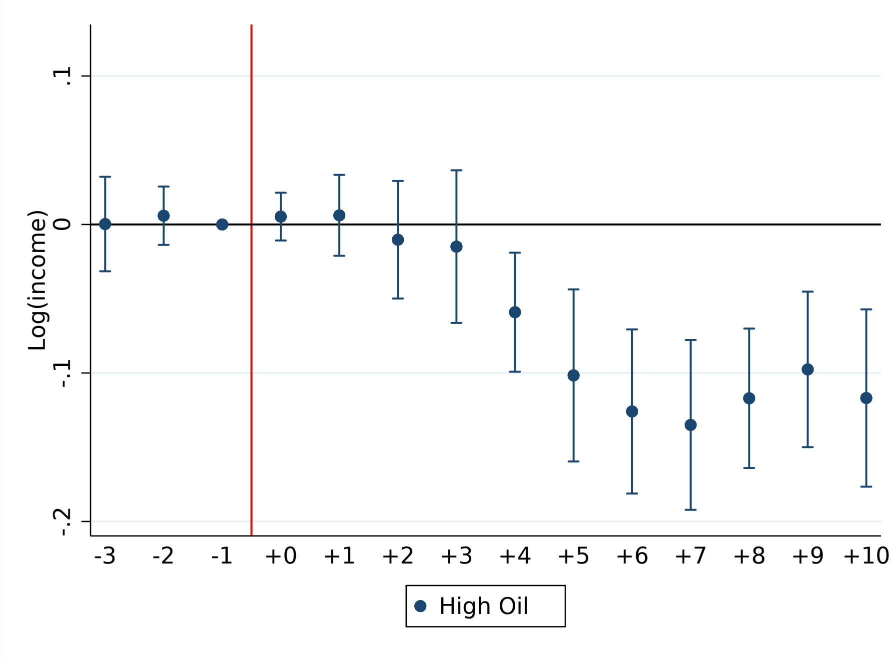

Figure S9: Estimated Effects of the Oil Discovery on Log(Income) over Time by Gender Excluding Oslo. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on log of annual income by gender, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated in a low oil local labour market - excluding Oslo, the capital - in 1965 in equation (1). Vertical line corresponds to discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

Panel A. Men

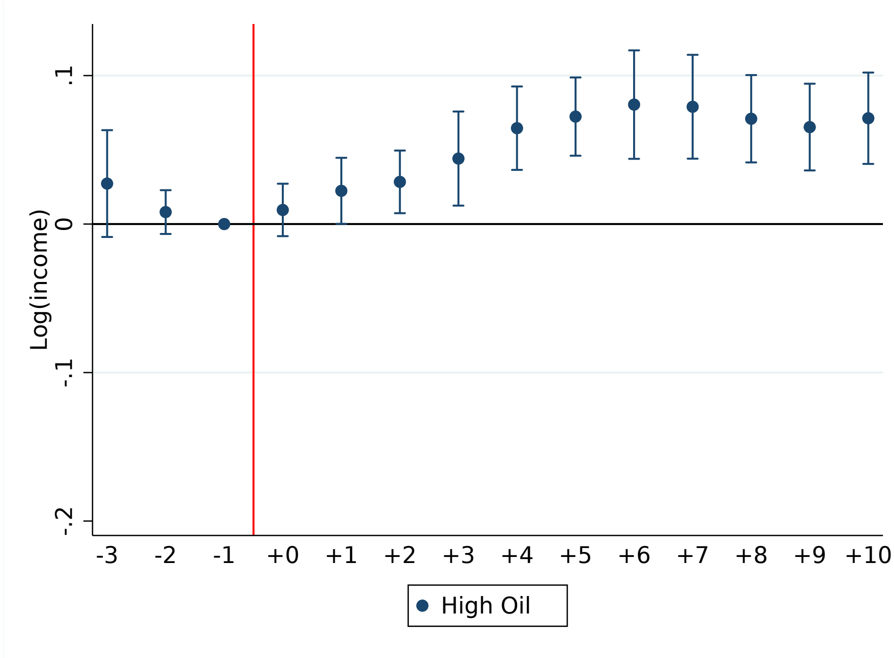

Panel B. Women

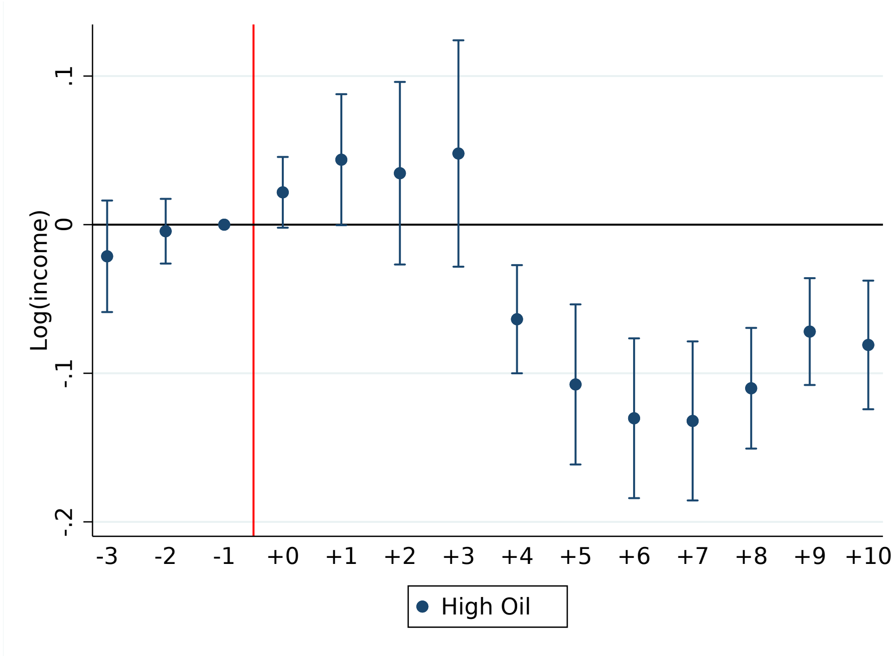

Figure S10: Estimated Effects of the Oil Discovery on Log(Income) over Time by Gender Excluding Northern Markets. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on log of annual income by gender, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated in a low oil local labour market in 1965, excluding Northern local labour markets, in equation (1). Vertical line corresponds to discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

Panel A. Men

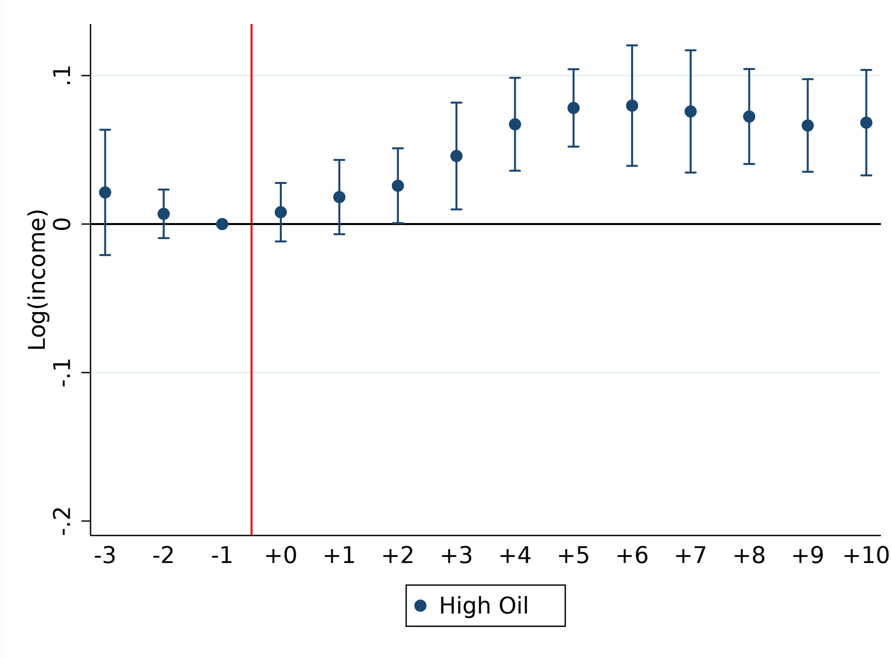

Panel B. Women

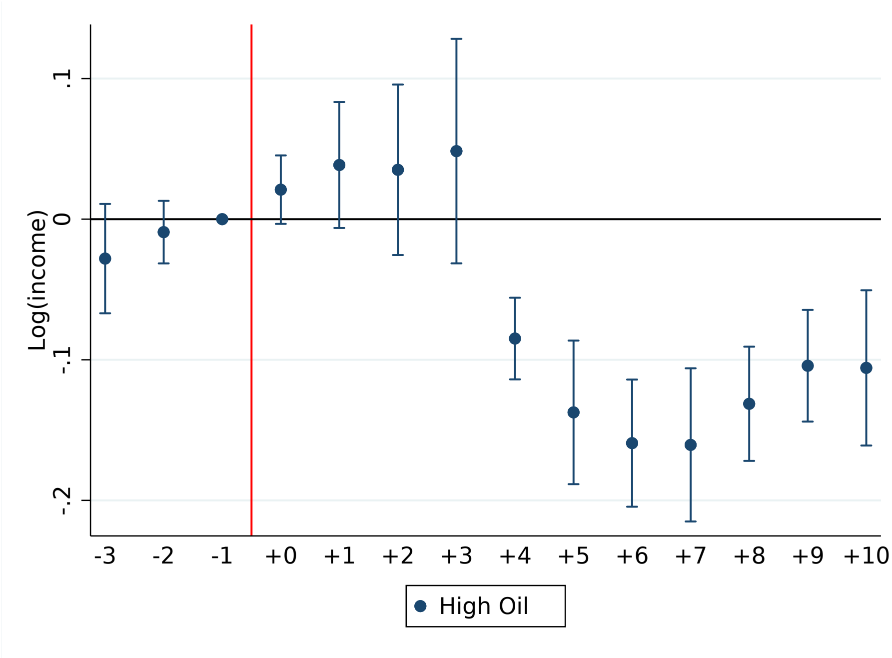

Figure S11: Estimated Effects of the Oil Discovery on Log(Income) over Time by Gender - Only Coastal Markets. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on log of annual income by gender, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated in a low oil local labour market in 1965 in equation (1). Estimation sample—high and low oil areas excluding any landlocked local labour market in both high- and low-oil. This isolates a sample of local labour markets which are, arguably, more comparable. Vertical line corresponds to discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

Panel A. Men

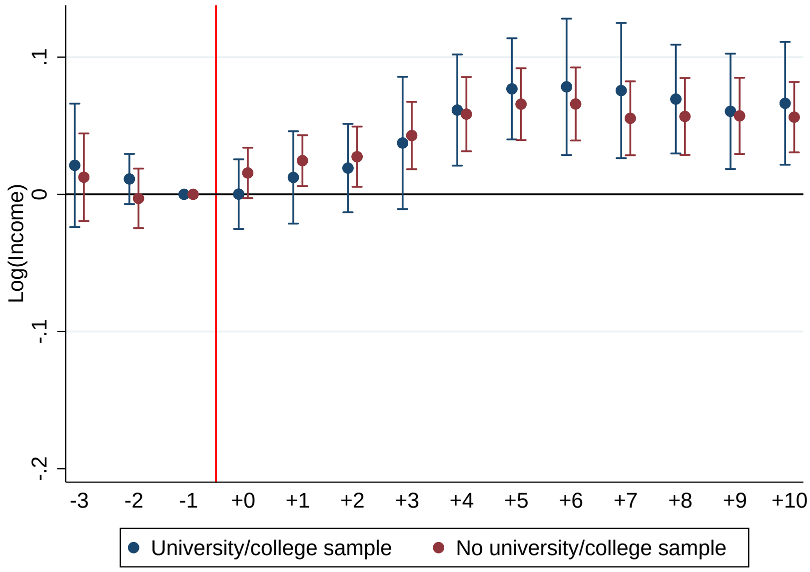

Panel B. Women

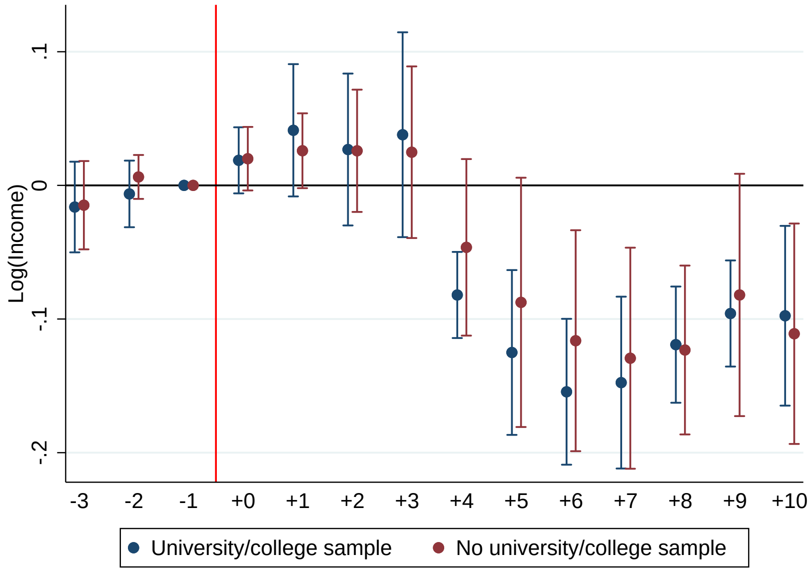

Figure S12: Estimated Effects of the Oil Discovery on Log(Income) over Time by Gender. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on annual log income by gender and by whether or not a local labour market has a university or college, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated in a low oil local labour market in 1965 in equation (1) from separate regressions by gender. The red vertical line corresponds to the discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero. 96,932 (51,300) men reside in local labour market with (without) a university or college respectively. 49,607 (21,007) women reside in local labour market with (without) a university or college, respectively.

D.0.5 Child care

In Figure S13, we show that there is no significant correlation between labour markets that experienced changes in available kindergartens and oil employment.

Figure S13: The correlation between Kindergarten Expansion and the Presence of Oil. Figure plots the correlation between the fraction of local labour market employment in oil in 1980 and the change in the kindergarten ("barnehage") enrolment rate between 1976–1979 as defined in Havnes and Mogstad (2011). Each point corresponds to a local labour market, weighted by the size of the population in 1970. Linear fit line plotted, not statistically significant at conventional levels.

D.0.6 Religiosity

Figure S14 plots the geographic distribution of the fraction of the population who is a member of the state church or any other church (panels (a) and (b) respectively). There is strong membership in the state church (which generally was the main religion for most Norwegians at the time) in the center of Norway, and to a lesser extent in the south-west, while there is stronger membership in any other church in the Southern parts of Norway. Separating state and any other church membership allows differences in beliefs to be accounted for: indeed, the south of Norway has much higher membership in other churches besides the state Church of Norway.

While there is a geographic dispersion of religiosity, note that even less religious places are still quite religious.

Figure S15 plots the correlation between the fraction of the population who is member of the state Church of Norway in 1960 and the fraction of employment in oil in 1980, while Figure S16 presents the same plot with membership in any other church. There is no significant relationship between the fraction of the people in a local labour market who are members of the state church or any other church and the fraction of employment in oil 20 years later. The estimated coefficients from both regression lines are small in magnitude and statistically insignificant. Furthermore, all our results are robust to the exclusion of the top 25% most religious areas in 1960. [30]

Figure S14: Geographic Distribution of Church Membership in 1960

(a) Membership in State Church

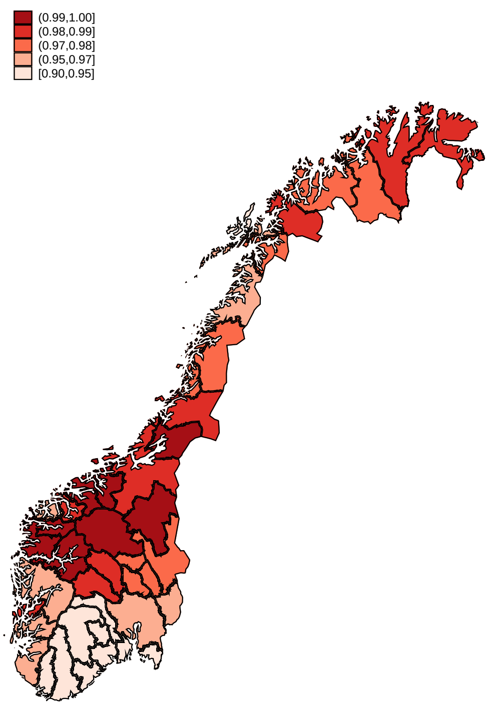

(b) Membership in Any Other Church

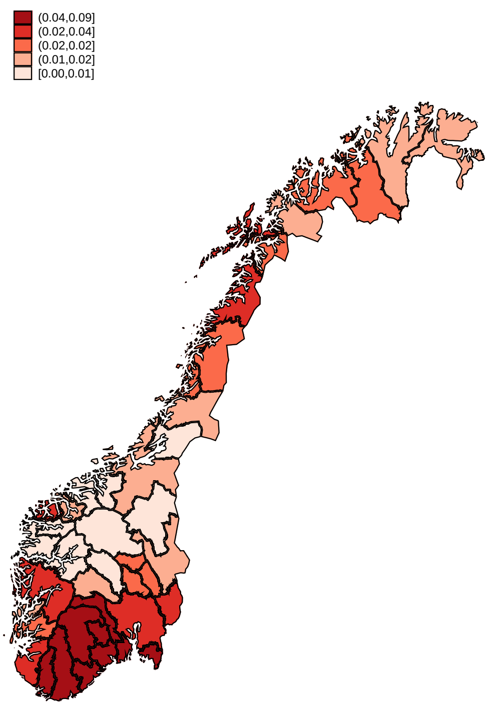

Note: Religiosity proxied as the fraction of residents who are a member of the state church (panel a) and any other church (panel b) in local labour markets in 1960.

[30]Results are available upon request.

Figure S15: The insignificant correlation between Religiosity and the Presence of Oil.

Note: Religiosity proxied as the fraction of residents who are a member of the state church in local labour markets in 1960. Each point represents a local labour market, weighted by the size of the population in 1960.

Figure S16: The insignificant correlation between Religiosity and the Presence of Oil.

Note: Religiosity measured as the fraction of residents who are a member of any other church in local labour markets in 1960. Each point represents a local labour market, weighted by the size of the population in 1960.

D.0.7 Time Horizon

Our results are robust to extending the time horizon to include the peak of the oil boom in the early 1980s, a time of high oil prices, and the oil ‘bust’ in the middle of that decade (in Figure S17). Up to 20 years after the boom, income gains for men in high local labour markets are persistent and stable, while for women income losses are persistent, though less severe over time.

Panel A. Men

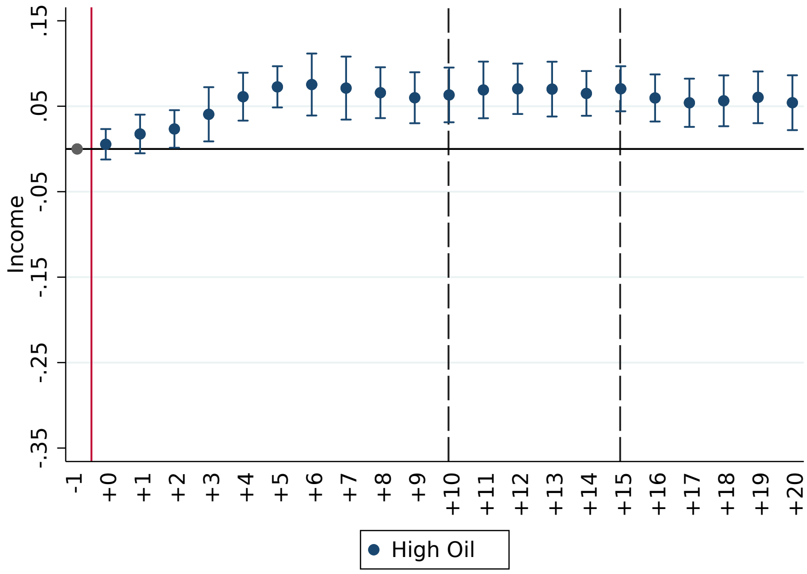

Panel B. Women

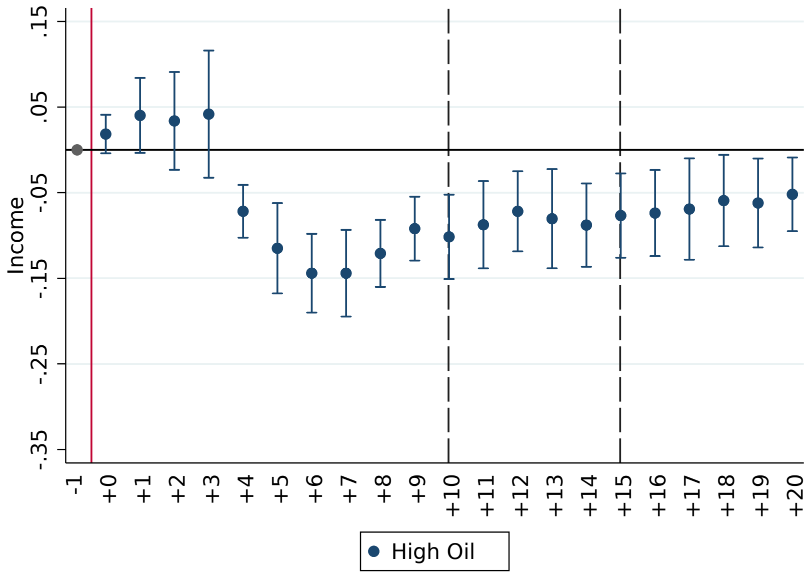

Figure S17: Estimated Effects of the Oil Discovery on Log(Income) over Time. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on annual log income by gender, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated in a low oil local labour market in 1965 in equation (1) from separate regressions by gender. The red vertical line represents the discover of oil in 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero. The dotted vertical line at +10 corresponds to the peak of global oil price in 1980, while that at +15 corresponds to the trough of oil price in 1985.

E Additional Estimation Results - Census

E.1 Income

Table S4: Census Results - Income

	Male		Female	
	(1) Income High Oil	(2) Income Middle Oil	(3) Income High Oil	(4) Income Middle Oil
Oil LLM=1	-13721***	-17965***	239	5951***
	(2288)	(1380)	(1134)	(1037)
Post=1	102581***	102581***	40628***	40628***
	(2527)	(2531)	(841)	(843)
Oil LLM=1 × Post=1	12083**	5979**	-6274***	-4170*
	(4567)	(2892)	(2272)	(2079)
Initial LLM FE	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Observations	445244	458264	420036	436280
R^2	0.279	0.279	0.083	0.081

Notes: Table reports the difference-in-differences estimates of equation (2) by gender, using a sample that includes birth cohorts between 1940-1949 (aged 21-30 in 1970). Dependent variable is income measured in NOK, inflation adjusted to year 1998. Treatment is high (columns 1 and 3) or middle (columns 2 and 4) exposure to oil based on 1965 residence. Standard errors clustered at the initial local labour market level are reported in parentheses. ***$\rho < 0.01$, **$\rho < 0.05$, *$\rho < 0.10$.

E.2 Additional Hours Worked Bins

Table S5: Census Results - Employed 1000–1299 hours

Sample:	Strongly Attached 67–69		Weakly Attached 67–69	
	(1) 1000–1299 Hours Worked Male	(2) 1000–1299 Hours Worked Female	(3) 1000–1299 Hours Worked Male	(4) 1000–1299 Hours Worked Female
Oil LLM=1	0.020***	0.035***	0.033***	0.014***
	(0.004)	(0.002)	(0.005)	(0.003)
Post=1	-0.067***	0.008***	-0.145***	-0.008**
	(0.005)	(0.002)	(0.003)	(0.003)
Oil LLM=1 × Post=1	0.003	-0.010*	0.003	0.002
	(0.008)	(0.005)	(0.010)	(0.004)
Initial LLM FE	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Observations	293879	112497	113303	136652
R^2	0.031	0.001	0.063	0.002
Mean Dep. Variable 1970	0.105	0.135	0.192	0.167

Notes: Table reports the difference-in-differences estimates of equation (2) using a sample that includes birth cohorts between 1940-1949 (aged 21-30 in 1970). Dependent variable is a dummy variable indicating employment of 1000–1299 hours annually. Sample is restricted to those with any employment in 1970 and 1980. Treatment is high exposure to oil based on 1965 residence. 'Strongly attached' to LF defined as having positive income in all years 1967–1969, while 'weakly attached' is having zero income in at least one year from 1967–1969. Standard errors clustered at the initial local labour market level are reported in parentheses. ***$\rho < 0.01$, **$\rho < 0.05$, *$\rho < 0.10$.

Table S6: Census Results - Employed 500–999 hours

Sample:	Strongly Attached 67–69		Weakly Attached 67–69	
	(1) 500–999 Hours Worked Male	(2) 500–999 Hours Female	(3) 500–999 Hours Male	(4) 500–999 Hours Female
Oil LLM=1	0.004***	-0.022***	-0.009***	-0.024***
	(0.001)	(0.005)	(0.003)	(0.007)
Post=1	-0.000	0.143***	-0.063***	0.049***
	(0.001)	(0.006)	(0.005)	(0.004)
Oil LLM=1 × Post=1	-0.004***	0.015	-0.000	0.008
	(0.001)	(0.012)	(0.006)	(0.011)
Initial LLM FE	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Observations	293879	112497	113303	136652
R^2	0.003	0.043	0.019	0.008
Mean Dep. Variable 1970	0.025	0.079	0.097	0.203

Notes: Table reports the difference-in-differences estimates of equation (2) using a sample that includes birth cohorts between 1940-1949 (aged 21-30 in 1970). Dependent variable is a dummy variable indicating employment of 500–999 hours annually. Sample is restricted to those with any employment in 1970 and 1980. Treatment is high exposure to oil based on 1965 residence. 'Strongly attached' to LF defined as having positive income in all years 1967–1969, while 'weakly attached' is having zero income in at least one year from 1967–1969. Standard errors clustered at the initial local labour market level are reported in parentheses. ***$\rho < 0.01$, **$\rho < 0.05$, *$\rho < 0.10$.

Table S7: Census Results - Employed 100–499 hours

Sample:	Strongly Attached 67–69		Weakly Attached 67–69	
	(1) 100–499 Hours Worked Male	(2) 100–499 Hours Worked Female	(3) 100–499 Hours Worked Male	(4) 100–499 Hours Worked Female
Oil LLM=1	0.005***	-0.024***	0.000	-0.032**
	(0.001)	(0.003)	(0.008)	(0.014)
Post=1	0.004***	0.124***	-0.095***	-0.010
	(0.001)	(0.003)	(0.012)	(0.015)
Oil LLM=1 × Post=1	-0.005***	0.040***	-0.004	0.041*
	(0.002)	(0.008)	(0.015)	(0.021)
Initial LLM FE	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Observations	293879	112497	113303	136652
R^2	0.002	0.047	0.039	0.002
Mean Dep. Variable 1970	0.017	0.053	0.120	0.208

Notes: Table reports the difference-in-differences estimates of equation (2) using a sample that includes birth cohorts between 1940-1949 (aged 21-30 in 1970). Dependent variable is a dummy variable indicating employment of 100–499 hours annually. Sample is restricted to those with any employment in 1970 and 1980. Treatment is high exposure to oil based on 1965 residence. 'Strongly attached' to LF defined as having positive income in all years during 1967–1969, while 'weakly attached' is having zero income in at least one year from 1967–1969. Standard errors clustered at the initial local labour market level are reported in parentheses. ***$\rho < 0.01$, **$\rho < 0.05$, *$\rho < 0.10$.

F Distribution of Occupations in 1970

Panel A. Men

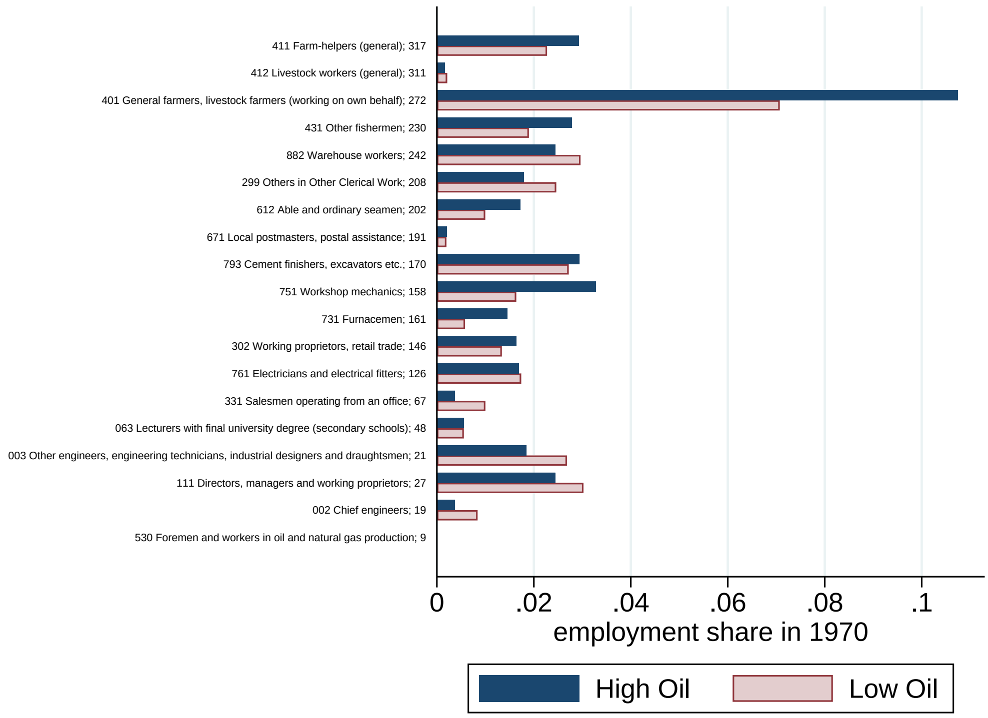

Panel B. Women

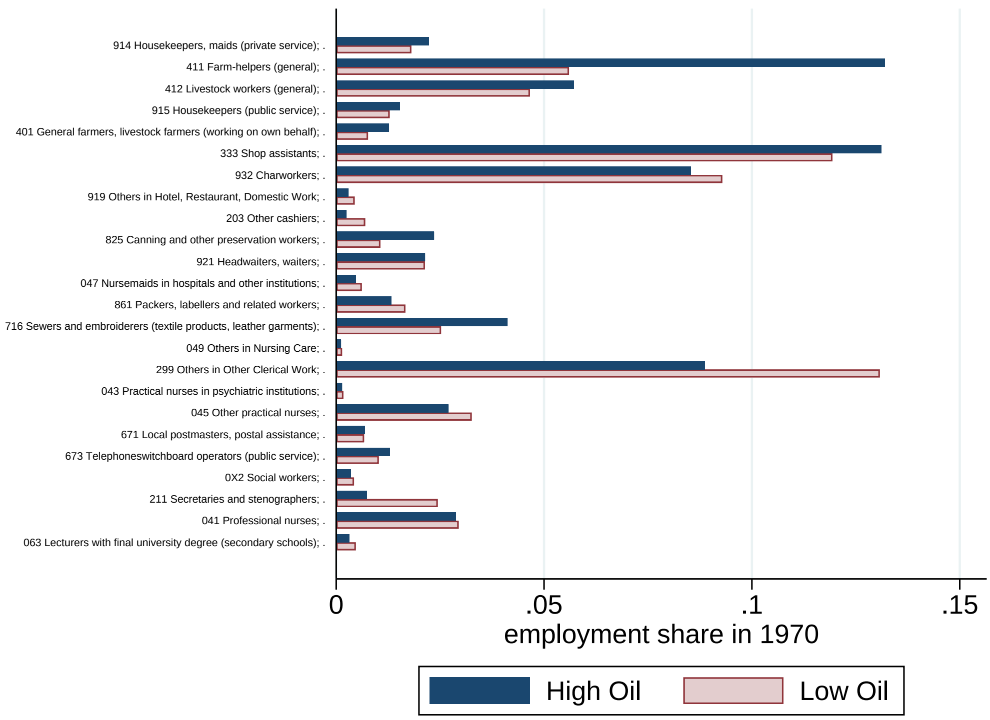

Figure S18: Employment Shares in 1970. The figure plots the employment shares in high versus low oil local labour markets in 1970 among the set of occupations in Figure 7. The number at the end of each occupation denotes the national rank of the occupation in terms of average income earned in 1980. Occupations classified according to the International Classification of Occupation (ISCO) system. Sample of workers in high/low employed in an occupation in a given year.

G Fertility

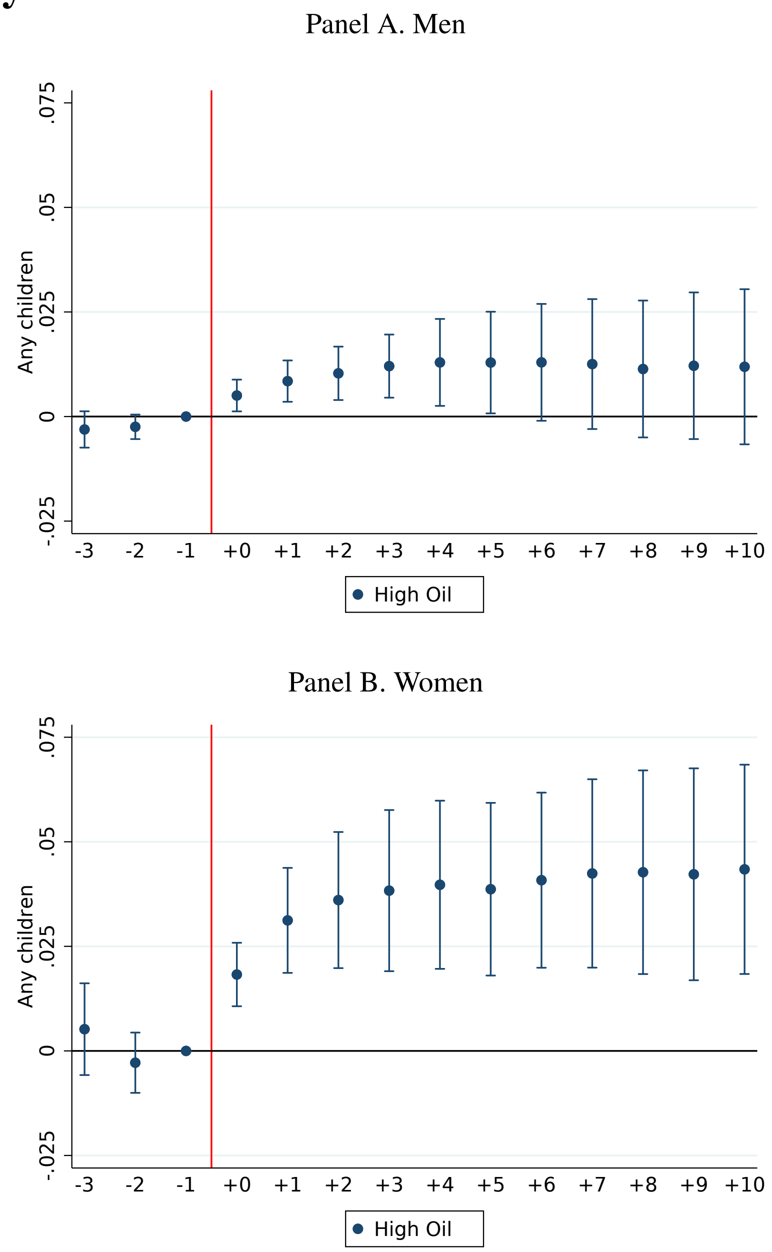

Figure S19: Estimated Effects of the Oil Discovery on Having Any Children over Time. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on a dummy variable indicating whether the person has any children, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated ones in low oil local labour markets in 1965 in equation (1). Vertical line corresponds to discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

Panel A. Men

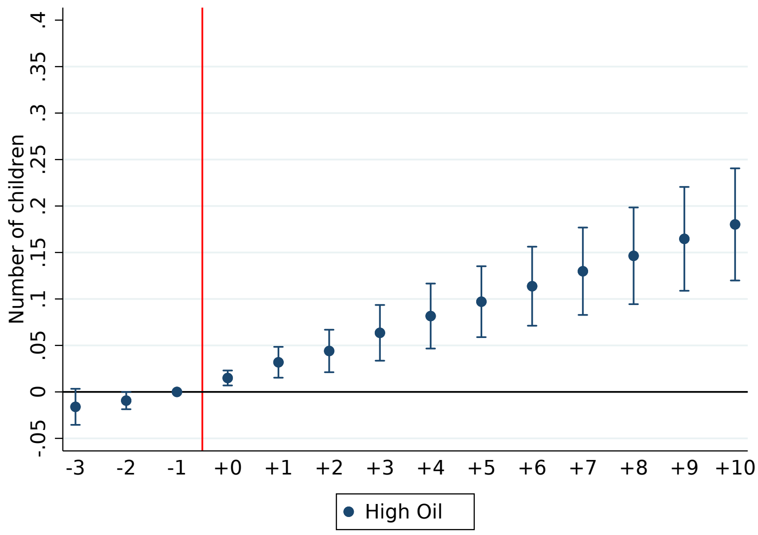

Panel B. Women

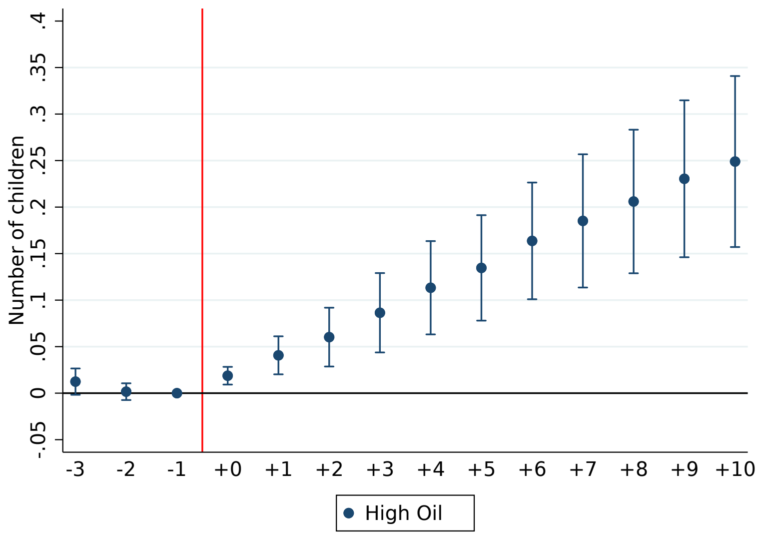

Figure S20: Estimated Effects of the Oil Discovery on the Number of Children over Time. Birth cohorts: 1940-1949. The figure plots the estimated impact of the oil discovery (δ_k) on the number of children that individuals had before and after the oil shock, comparing treated individuals who resided in a high oil local labour market in 1965 to untreated ones in low oil local labour markets in 1965 in equation (1). Vertical line corresponds to discovery of oil at the end of 1969. 95% confidence interval reported. Reference period -1 (1969) is normalized to zero.

H Transition Between Labour Statuses

We examine the transition probabilities of women who had an occupation in 1970, considering where they end up in 1980. First, we consider women with any occupation, and afterwards women with full time employment. Those who have an occupation in 1970, ten years later can either: (a) have the same occupation, (b) have a new occupation, or (c) have no occupation. Those in full time unemployment could also have moved to part-time employment.

We find that the effect of oil on women's employment largely operates through an "exit from the labour force" channel, as women are more likely to have become unemployed in 1980 than low oil women. Women in high oil regions are also somehow less likely to move to a new occupation, but occupational mobility does not show any difference for women who have an occupation both in 1970 and 1980. Also, movements towards part-time jobs are not a major force behind the changes experienced by female workers, as the decline in full-time employment in high oil is met nearly one-to-one with the increase in unemployment.

Figure S21 shows that high oil women are more likely to have become unemployed in 1980 than low oil women. They also are somehow less likely to be in a new occupation (40% in high oil compared to 45% in low oil), so there is an "out of the labour force" channel as well as possibly a "lower occupational mobility" channel. However, if we focus on women who have an occupation both in 1970 and 1980, then the decrease in occupational mobility is gone. Thus, this decline in occupational mobility is observed because there is an increase in the probability of having no occupation in high oil in 1980, suggesting that changes in occupations are not a major channel, while dropping out of the labour force is more important.

Next, we repeat the above exercise, but conditioning on full-time employment in 1970 (defined as 1300 hours or more annually). By 1980, those women who are employed full-time in 1970 can either: (a) still be employed full-time, (b) be employed part-time (less than

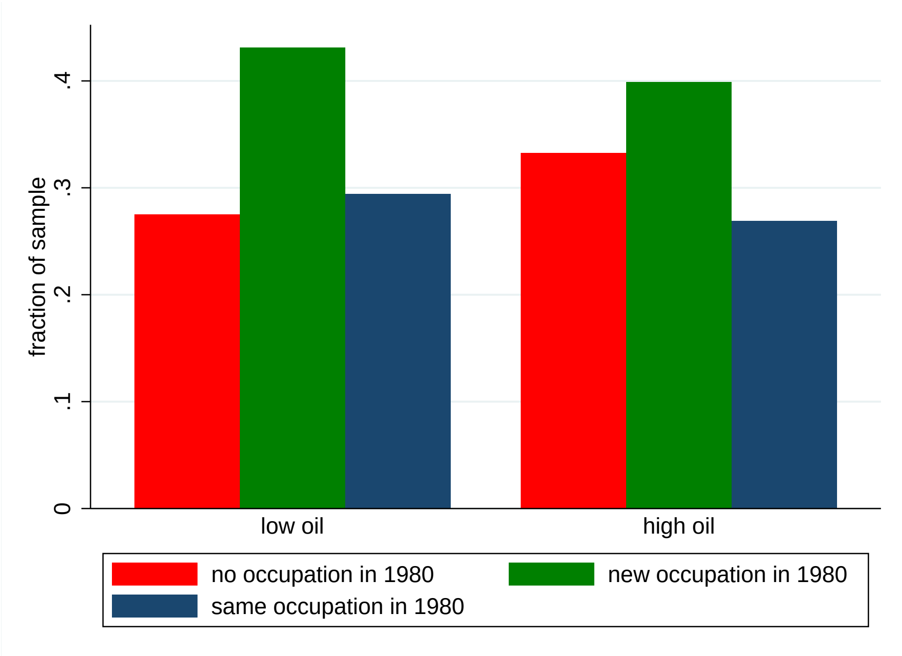

Figure S21: Figure graphs the probability of having no occupation in 1980, a new occupation in 1980, or the same occupation in 1980 among the sample of women who have an occupation in 1970, separately by high and low oil.

1300 hours annually), or (c) be unemployed. We present the results in Fig. S22. Generally, we see a decline in full-time in high oil which is met nearly one-to-one with the increase in unemployment. Once again, it seems that the main mechanism is due to working women leaving the labour force.

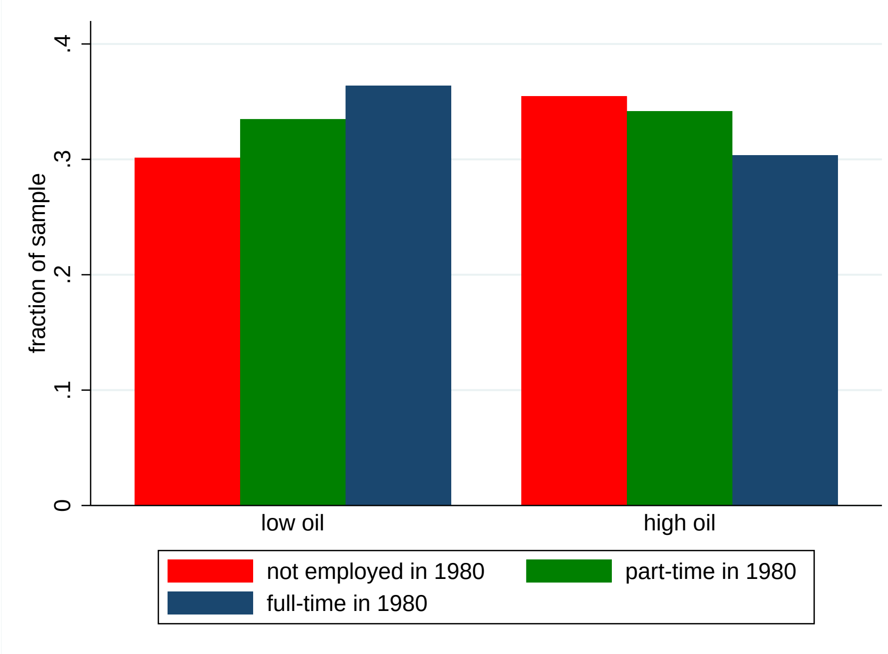

Figure S22: Figure graphs the probability of being employed full-time in 1980, part-time in 1980, or non employed in 1980 among the sample of women who have an occupation in 1970 and are employed full-time, separately by high and low oil.

I Long-Run Estimation and Mechanisms

In order to test if the effect of the oil shock has a different effects on later cohorts of workers who were children at the time of the boom, we estimate the impacts of the oil discovery on lifetime income as follows:

$$LY_{icab} = \sum_{a=[4\text{-}6]}^{[19\text{-}21]} (Oil_c \times P_{ia})\alpha_a + \psi_c + \tau_b + u_{icb} \tag{3}$$

where LY_{icab} denotes the lifetime income of individual i who resides in local labour market c prior to 1969, of age group a at the time of the oil shock (where age bins are $a = 4\text{-}6, 7\text{-}9, \ldots, 19\text{-}21$) and in birth cohort b, which includes all birth cohorts from 1949 to 1966. Oil_c is an indicator variable denoting whether an individual initially resides in a high oil local labour market. P_{ia} denotes the age bin of individual i in 1970. The omitted reference group consists of individuals aged 19-21 in 1970. ψ_c denotes the time-invariant differences across local labour markets; τ_b represents differences across birth cohorts. Standard errors are clustered at the initial local labour market level.

Table S8 presents the coefficients of the interactions for men and women (columns 1 and 2 respectively). Each coefficient α_a represents the difference in lifetime income between individuals of age group a in high versus low oil compared to the difference in lifetime income between individuals aged 19–21 in high relative to low oil markets. Women aged 19–21 in high oil have an average lifetime income 9.5% *lower* than women in low oil. As such, a positive coefficient corresponds to a closing of the negative gap in lifetime income between high and low oil women aged 19–21. For men, those in high oil have an average lifetime income 5.5% *higher* than men in low oil and a positive coefficient corresponds to a widening of the positive gap in lifetime income between high and low oil men aged 19–21.

Table S8: Impact of Oil Shock on Log(Lifetime Income) by Age in 1970

Sample:	Male	Female
	(1) Log(Lifetime Income) High Oil	(2) Log(Lifetime Income) High Oil
19-21	ref.	ref.
16-18	-0.008	0.215***
	(0.007)	(0.011)
13-15	0.015**	0.328***
	(0.005)	(0.016)
10-12	0.053***	0.444***
	(0.006)	(0.022)
7-9	0.135***	0.578***
	(0.013)	(0.017)
4-6	0.133***	0.625***
	(0.014)	(0.024)
Oil × 16-18	0.012	0.043**
	(0.008)	(0.017)
Oil × 13-15	0.024*	0.042**
	(0.012)	(0.017)
Oil × 10-12	0.012	0.060**
	(0.018)	(0.025)
Oil × 7-9	0.004	0.066**
	(0.020)	(0.028)
Oil × 4-6	0.020	0.072***
	(0.027)	(0.025)
Initial LLM FE	Yes	Yes
Cohort FE	Yes	Yes
Observations	430140	407651
R^2	0.014	0.045

Notes: Table reports the difference-in-differences estimates of equation (3) using a sample that includes birth cohorts between 1949-1966 (aged 4-21 in 1970). Dependent variable is lifetime income, defined as cumulative income earned between ages 30-45. Income measured in NOK, inflation adjusted to year 1998. Reference group is those aged 19-21. Men aged 19–21 in high oil have an average lifetime income 5.5% greater than men in low oil while women aged 19–21 in high oil have an average lifetime income 9.5% lower than women in low oil. Treatment is high exposure to oil based on 1965 residence. Standard errors clustered at the initial local labour market level are reported in parentheses. ***$\rho < 0.01$, **$\rho < 0.05$, *$\rho < 0.10$.

Table S8 confirms the patterns of Figure 8, namely that the age at the time of the oil discovery matters for the lifetime income of women.

Many potential factors could explain how younger cohorts of women close the lifetime income gap with their peers in non-oil regions. Labour supply among younger women may increase relative to older women, as those who were children at the time of the economic shock had more time to consider and adapt their labour supply choices. In addition, younger women may alter their levels of education in response to the local demand shock, as documented by Black et al. (2005) and Cascio and Narayan (2015), who find a causal link between natural resource booms and changes in education. We present suggestive evidence that the reduction in women's earnings losses for the next generation of female workers is driven by changes in labour market participation. While many other economic and social factors may adjust simultaneously over such a long time horizon, it remains useful to examine how the labour supply and education decisions of the next generation depend on the age at which women and men are exposed to the economic shock and may play a role in explaining the recovery of women's income.

First, we show that different rates of long-term migration between women in high- and low-oil are similar across different ages of exposure, suggesting migration is not an important channel (Appendix I.1). Second, making use of employment status available annually from 1986–2017, we show that full-time employment measured at age 44 among older women in high-oil lags behind low-oil (Appendix I.2). However, similar to the catch up observed for earnings, full-time employment of those exposed to oil from 4–6 in high-oil overtakes women in low-oil, suggesting a role of labour supply decisions over time. Third, we find no evidence that younger cohorts of women modify their education decisions with respect to high school or post-secondary education in response to the local shock (Appendix I.3.).

I.1 Migration

First, important differences in our results by age at exposure suggest that the initial disadvantage among working age women disappears over time as the next generation has more flexibility to adapt formative decisions in response to the local economic boom. One channel through which lifetime income improves among the youngest women could be due to different rates of migration. Wilson (2020), for example, documents substantial migration into oil regions following the fracking boom in the United States. As discussed previously, migration rates in the older generation does not differ substantially between high and low oil workers, but it is important to verify that this pattern also holds in the long run. Since the younger cohorts had more time to relocate to other labour markets in response to the economic boom, we need to ensure that the lifetime income gains we observe for high-oil women are not driven by women choosing to move, for example, to Oslo, the national as well as business capital. Since our definition of treatment is based on the area of residence in 1965, prior to the local boom, if we observe extensive mobility in the long run, this could be the main margin of adjustment for female workers. Figure S23 rules out mobility as a key explanatory factor. Women in the high-oil regions tend to relocate systematically less than those living in low-oil areas, and the difference in migration rates between the two regions is stable across the youngest cohorts.

I.2 Labour Supply

We have already showed that young working women significantly reduce their participation in full-time employment in the aftermath of the boom. If this trend was only temporary, it could explain the observed catching up of female income in high oil regions in the next generation. Figure S24 illustrates that, indeed, while most cohorts of men in high oil labour markets has a higher share of workers in full-time employment than their low-oil peers by the age of 44, only

Panel A. Men

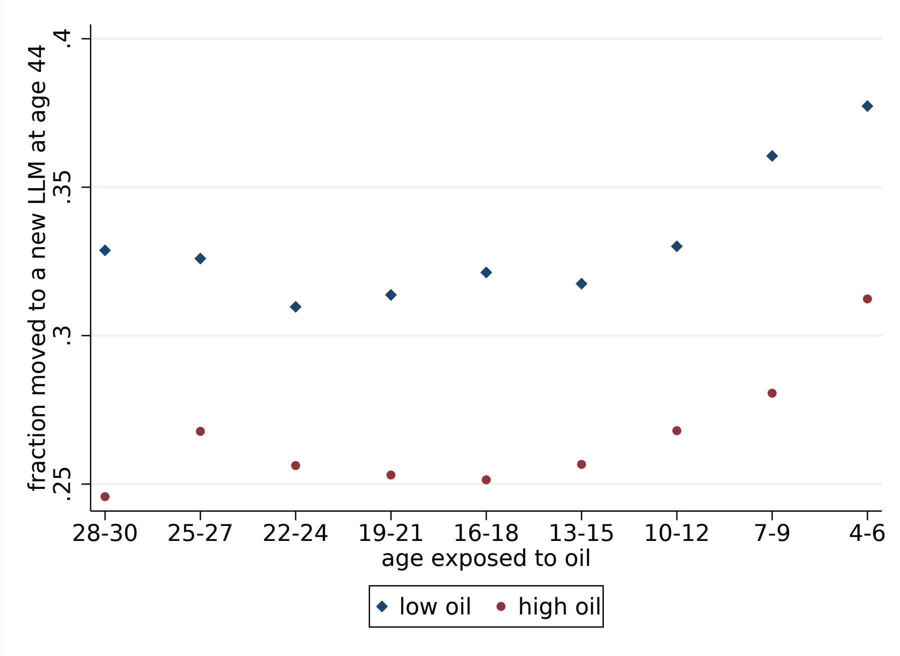

Panel B. Women

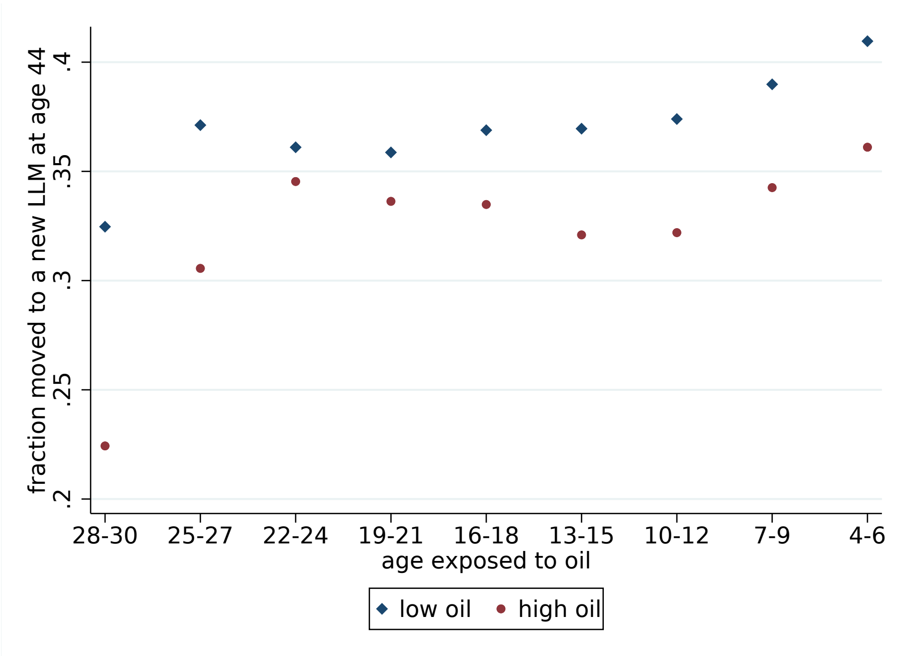

Figure S23: Share of migration to new local labour markets by cohorts and gender. The figure plots the share who have moved to a different local labour market, measured at age 44, of each age group by the age exposed to the local boom. High/low oil defined as residing in high/low oil local labour markets in 1965 (from birth for those born after 1965).

the youngest cohorts of women (aged 4–6 years old in 1969) has a larger fraction of workers in full-time employment than their low-oil counterparts. Consistent with the lack of income differences between young and older men in Table S8, the difference in full-time employment between men in high and low oil regions is stable between those aged 19–21 and those aged 4–6 at oil discovery. As the gap in full-time employment between women in the treated and control regions closes, so should the lifetime earnings gap, *ceteris paribus*. We then have a first indication that the catching-up in female lifetime earnings could come from an adjustment in full-time employment.

I.3 Education

Beyond a direct adjustment in full-time labour supply, it is also plausible that the local boom led to some long-term adjustments in education. The new labour market situation, with more high-paying jobs for men, but fewer ones for women, changed the opportunity cost of remaining in school - the foregone income from remaining outside of employment for an additional year. Thus, education and training could be one transmission channel of the oil shock that differs for the next generation of workers and could explain the catching-up of women's earnings in the high-oil regions. Since it takes a long time to train as petroleum engineers and other similar high-paying jobs that grew in demand after the economic boom, the next group of workers could benefit from the option to adjust educational choices after learning about the oil discovery.

Moreover, if children experienced a higher household income because of the economic boom, they could afford more education, and this could be more relevant for girls who otherwise would not invest in those types of training. Table S10 in the Appendix shows the interaction coefficients of the shock with two measures of education. Overall, we do not find evidence

Panel A. Men

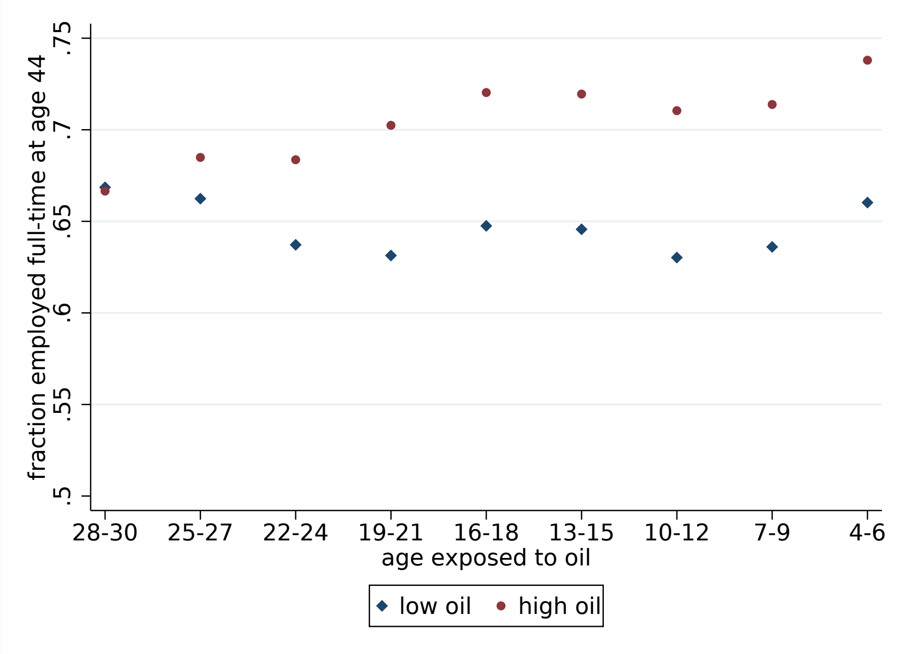

Panel B. Women

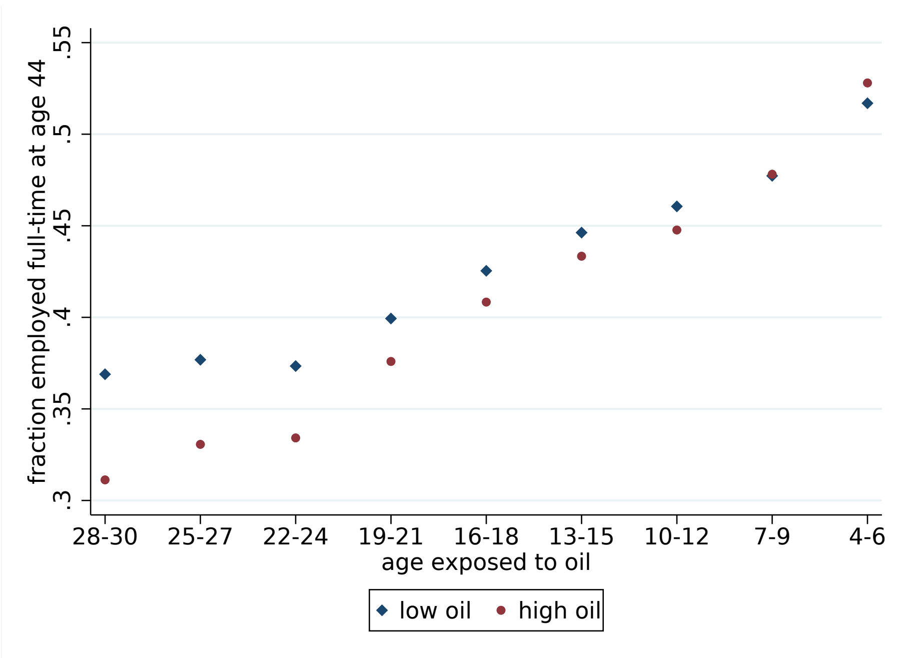

Figure S24: Share of full-time employment by cohorts and gender. The figure plots the share employed full-time, measured at age 44, of each age group. High/low exposure to oil is defined as residing in high/low oil local labour markets in 1965 (or from birth for those born after 1965). Employment information from annual employment data, where full-time employment is defined as employment of at least 30 hours a week.

that girls modify their educational trajectories in response to the local boom. Columns (1) and (3) examine the effect of the shock on high school completion for men and women of different age bins, respectively, while columns (2) and (4) examine the effect on post-secondary completion. In general, there are no substantial changes in these measures of education. Teenage boys aged 16-18 are slightly less likely to complete post-secondary education in the boom regions, while the youngest cohort of 4-6 year olds are slightly more likely to finish high school, but there are no statistically significant changes for women. Thus, education seems to not be a major factor behind the catch up of younger women in high oil over time.

Table S9: Impact of Oil Shock on Years of Education of Young Workers

	Male	Female
	(1) Years of Education High Oil	(2) Years of Education High Oil
Oil LLM=1	0.083***	0.062***
	(0.008)	(0.007)
Post=1	0.693***	0.333***
	(0.012)	(0.008)
Oil LLM=1 × Post=1	0.041**	-0.012
	(0.016)	(0.013)
Initial LLM FE	Yes	Yes
Cohort FE	Yes	Yes
Observations	443816	419114
R^2	0.036	0.027
Mean Dep. Variable	10.274	9.737

Notes: Table reports the difference-in-differences estimates of equation (2) using a sample that includes birth cohorts between 1940-1949 (aged 21-30 in 1970). Dependent variable is completed years of education. Treatment is high exposure to oil based on 1965 residence. Standard errors clustered at the initial local labour market level are reported in parentheses. ***$\rho < 0.01$, **$\rho < 0.05$, *$\rho < 0.10$.

Table S10: Impact of Oil Shock on Lifetime Education of Next Generation (by Age in 1970)

	Male		Female	
	(1) Completed HS High Oil	(2) Completed Post-Sec High Oil	(3) Completed HS High Oil	(4) Completed Post-Sec High Oil
19-21	ref.	ref.	ref.	ref.
16-18	0.035***	0.033***	0.084***	0.065***
	(0.009)	(0.005)	(0.004)	(0.004)
13-15	0.088***	0.029***	0.146***	0.085***
	(0.009)	(0.005)	(0.006)	(0.005)
10-12	0.146***	0.018***	0.229***	0.111***
	(0.015)	(0.004)	(0.016)	(0.007)
7-9	0.193***	0.036***	0.315***	0.140***
	(0.018)	(0.005)	(0.016)	(0.010)
4-6	0.223***	0.062***	0.401***	0.185***
	(0.022)	(0.008)	(0.026)	(0.018)
Oil × 16-18	0.004	-0.010**	-0.003	-0.002
	(0.007)	(0.005)	(0.008)	(0.008)
Oil × 13-15	0.015	-0.002	-0.001	-0.003
	(0.009)	(0.008)	(0.009)	(0.010)
Oil × 10-12	0.025	0.000	0.000	-0.006
	(0.015)	(0.008)	(0.018)	(0.014)
Oil × 7-9	0.025*	-0.008	0.001	-0.011
	(0.014)	(0.012)	(0.023)	(0.020)
Oil × 4-6	0.050**	-0.000	0.013	0.001
	(0.020)	(0.014)	(0.034)	(0.029)
Initial LLM FE	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Observations	440881	440881	420067	420067
R^2	0.034	0.006	0.072	0.017

Notes: Table reports the difference-in-differences estimates of equation (3) using a sample of next-generation workers that includes birth cohorts between 1949 -1966 (aged 4-19 in 1970, at the start of the boom). Dependent variable is a dummy variable indicating if a level of education (High school or Post-secondary education) was attained by the age 40. Reference group comprises those aged 19-21. Treatment is high exposure to oil based on 1965 residence. Standard errors clustered at the initial local labour market level are reported in parentheses. ***$\rho < 0.01$, **$\rho < 0.05$, *$\rho < 0.10$.