Household Responses to Export Prices - Evidence from an Oil Exporting Country*

Andreas Fagereng† Magnus A. H. Gulbrandsen‡ Gisle J. Natvik§

February 14, 2022

Abstract

How do shocks to a country's export prices transmit to its households' incomes and affect their behavior? With fine-grained Norwegian micro data, we estimate how households in a small open economy are affected by and respond to variation in the price of their country's main export article, oil. Following an oil price hike, households across sectors experience a hump-shaped positive income response. This response is stronger for households with earnings from more oil-exposed sectors, such as the oil extraction industry. We rely on this differential exposure to estimate (i) households' intertemporal propensities to spend out of future income changes, (ii) spousal earnings responses, and (iii) general equilibrium spillover effects to households outside the oil sector, all three central mechanisms in the burgeoning literature that aims to reconcile models of macroeconomic shock propagation with microeconomic evidence. Our main findings are as follows. (i) While household income responds with a lag, expenditure responds on impact, implying an intertemporal propensity to spend out of income that is frontloaded and consistent with anticipation effects in consumption. (ii) The earnings of oil-workers' spouses drop after oil price hikes, consistent with substantial inter-household risk-sharing along the labor margin. (iii) The higher is the share of oil-sector households living in a region, the more are households employed outside the oil sector affected by oil price changes, consistent with sizable general equilibrium spillover effects from the oil sector to non-oil sector workers. Additional evidence across industries and occupations suggests these spillovers mainly result from higher oil-sector labor demand, not Keynesian-style expenditure multiplier effects.

*This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. We thank Christian Brinch, Peter Fredriksson, Martin Holm, Espen Moen, Roine Vestman, and various seminar participants for useful comments and suggestions.

†BI Norwegian Business School, email: andreas.fagereng@bi.no. Also affiliated with Statistics Norway.

‡Norges Bank and BI Norwegian Business School, email: magnus.a.gulbrandsen@gmail.com.

§BI Norwegian Business School, email: gisle.j.natvik@bi.no. Also affiliated with Statistics Norway.

1 Introduction

In an open economy, households are potentially exposed to variation in the globally determined prices of the goods and services that their employers export. Traditional expositions in macroeconomics highlight how improvements in a country's terms of trade trickle down to household income, lifts consumption expenditure and possibly unleashes further domestic propagation through Keynesian multiplier effects (Fleming, 1962; Dornbusch and Fischer, 2017). Recent macroeconomic research formulates the potential channels more precisely, emphasizing how households optimally respond to export price movements and how these responses contribute to the macroeconomic consequences of international shocks for small open economies (Obstfeld and Rogoff, 1995; Galí and Monacelli, 2005; Auclert, Rognlie, Souchier, and Straub, 2021). However, the theoretical narratives of household-level propagation mechanisms in open economies have to a limited degree been guided by evidence on how export price variation actually affects households. In this paper, we seek to provide such evidence.

We investigate how households in Norway are affected by and respond to changes in the country's main export article, oil.[1] With detailed and precise administrative data covering the entire population over nearly two decades, we estimate households' income- and expenditure responses, conditional on their labor market and regional exposure to oil activity. We focus on three dimensions that are central in recent attempts to build models of macroeconomic shock propagation that are consistent with micro-level facts.[2] First, we estimate how households' expenditure responds to the future income changes that oil price hikes induce. Second, we look within the household and estimate the extent to which spousal earnings provide insurance of the idiosyncratic income risk that oil price fluctuations implies for oil workers. Third, we look across municipalities to explore the extent to which oil price changes spill over to households employed outside the oil sector through local general equilibrium (GE) effects. In addition, we exploit variation across industries within the non-oil sector and across different occupations to assess whether our estimated GE effects come from a "neoclassical" labor demand channel or a Keynesian aggregate demand channel.

Our empirical strategy builds on the presupposition that oil price fluctuations are an exogenous source of variation in income.[3] Figure 1 displays our study's starting point. The

[1] Over our sample period, 1996 to 2014, oil sector exports were a source of massive revenues to Norway, accounting for about half of total Norwegian export income. In total, oil-related activities accounted for 15 to 25 percent of Norwegian GDP over these years. The economy has therefore been sensitive to oil price swings, as documented by empirical studies on macroeconomic time-series (see e.g. Bjørnland and Thorsrud (2016), Bjørnland, Thorsrud, and Torvik (2019) and Bergholt, Larsen, and Seneca (2019)).

[2] See for instance Kaplan, Moll, and Violante (2018), Auclert, Rognlie, and Straub (2018), Auclert, Rognlie, and Straub (2020b), and Bardóczy (2020).

[3] Other studies have previously used aggregate oil price fluctuations to identify various effects of changing household income at a more disaggregated level. Among the notable examples are Acemoglu, Finkelstein,

upper marked line is the estimated income response to a ten percent oil price hike among households where at least one member earns his or her labor income from the *oil sector*.[4] The dashed curve in the middle represents the income response among households with at least one earner in the *oil service sector*. Sector definitions here follow the ones used by Statistics Norway, Section 2 provides the details. The bottom solid curve displays the income response of households without any income from either of these oil-related sectors. A priori, we would expect that the oil-sector households are the most exposed to oil price changes, the non-oil households are the least exposed, and the oil-service households are somewhere in between. The figure is consistent with this prior. All households experience greater income for a while, but more so for the households with earnings from the oil sector.

Figure 1: Income responses to a ten percent change in the oil price

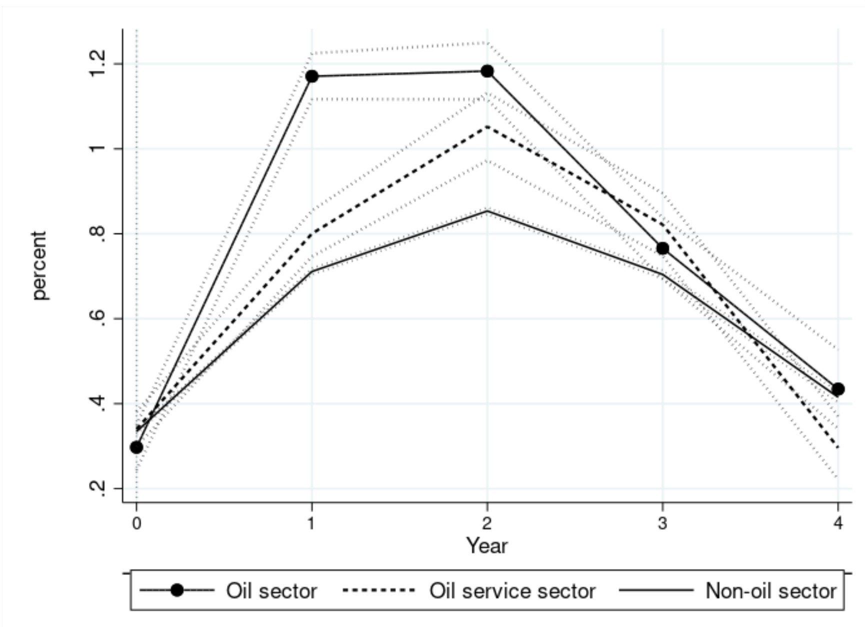

Of course, oil prices change for a variety of reasons and they are likely to coincide with many other macroeconomic events. It would therefore be naive to interpret the income responses in Figure 1 as the causal effects of a ten percent oil price hike. For this reason, our empirical strategy is to exploit the heterogeneity in exposure to oil production to obtain estimates that we interpret causally. With changes in the oil price and an indicator for whether households obtain income from either of the oil industries, we construct a shift-

and Notowidigdo (2013) who study effects of income on health expenditures, and Charles, Li, and Stephens Jr (2018) study how local labor market conditions affect disability take-up.

[4]In our sample period from 1996 to 2014 the mean yearly percent change in the oil price was 13 percent, and the standard deviation was 25 percent.

share type variable as our main explanatory variable. These are applied in local projections Jordà (2005), similar to the approach that Chodorow-Reich, Nenov, and Simsek (2021) use to estimate effects of stock prices in the US. Notably, we purge the micro responses of time-fixed effects which necessarily are perfectly correlated with oil price fluctuations. Our inference is thereby obtained by comparing households in the oil and oil-service sectors to households in the non-oil sector. In other words, our study focuses on differences in responses for oil and oil-service households versus non-oil households. One such difference is visualized in Figure 1. Moreover, because oil activity is geographically concentrated in Norway, we also control for region-time fixed effects, with regions defined by households' municipality of residence. These fixed effects control for not only nationwide conditions that covary with oil prices, but also for general equilibrium effects of oil price changes occurring at the municipal level. In subsequent analyses, we attempt to quantify how strong these effects are, by estimating how the share of oil households residing in a municipality affects non-oil households' responses to oil price changes.

Our first finding regards the relative timing of income and expenditure responses following an oil price change. Income responds in a hump-shaped manner. On impact, there is hardly any differential response of the oil versus the non-oil households. Then, a gap builds up before it closes again after four years, as one would expect if the labor market is competitive over time.[5] About 60 percent of the total five-year change in income arrives in year one after the price innovation, and most of the remainder arrives in year two.

In contrast to income, oil-sector households' expenditure increases immediately. It remains elevated for two more years, and thereafter drops. It even drops into negative territory, indicating that planned future expenses are shifted forward in time. The joint pattern of expenditure and income together imply distinct anticipation effects, assuming that the future income growth spurred by an oil price hike is predicted by households. Based on the time profile of income and expenditure responses, we compute the intertemporal propensities to spend out of (the sum of) future income changes. Our estimates suggest that out of the total 5-year income response, approximately 30 percent is spent by households in the impact year, approximately 75 percent is spent by the end of year one, and the rest is spent in year two after the price change.

On average across oil-sector households, the initial expenditure increase is primarily financed by debt, and only to a small extent by deposits. Thus, our estimates suggest that rather than running down current savings, oil households are able and willing to borrow (or cut back on amortization) against the future income hike, to finance greater spending immediately. In an extension of our main analysis, we show that this pattern does not differ

[5] In an extension to our main analysis using occupations data, we find that households employed in similar occupations as oil workers (e.g. engineers) experience a smaller wage differential that closes faster, again consistent with the implications of competitive markets from textbook economics.

to any notable extent between households with high or low initial liquidity, as measured by their deposit holdings or home equity.

Our second finding regards spousal income responses. As is well-known, a potentially important insurance mechanism available to two-adult households is family labor supply (Blundell, Pistaferri, and Saporta-Eksten (2016)). Therefore, we next turn to estimating how the incomes of oil worker spouses respond to oil price changes. To this end, we focus on households where only one member earns his or her income from the oil sector, and estimate the total income and earnings response of their spouses. Our estimates are negative, consistent with sizable spousal insurance of the idiosyncratic income risk induced by oil prices. A 10 percent oil price hike is followed by an accumulated 1.25 percent drop in spousal earnings over the next four years. In years one and two, the negative response of spousal earnings is approximately half that of the oil-sector workers. In year three, total oil-sector household income has approximately reverted to the same level as that of non-oil sector household income, but this reversion is due to a negative spousal earnings response that cancels out a positive oil-worker earnings response. The estimates indicate that the spousal income margin is an important source of household-level insurance.

Because oil-sector households tend to live in oil-intensive regions, their earnings after an oil price change are likely to be elevated by local general equilibrium effects. Similarly, the spousal earnings response is potentially mitigated. In our baseline results, we control for these by including municipality-time fixed effects in our regressions. If we instead use time fixed effects only, the estimated spousal earnings response is indeed slightly more negative than our baseline estimates suggest, consistent with local spillovers from oil activity to non-oil workers' earnings. One possible explanation is that higher local expenditure by oil-sector workers elevate the income of non-oil workers through Keynesian-type demand multiplier effects. An alternative explanation is that there exists a regional component in the labor demand from the oil sector itself, so that after an oil price hike the oil sector increases its demand for labor more in those localities where there already are more oil-sector workers. We refer to both these mechanisms together as general equilibrium effects, and we first attempt to quantify their joint influence before we finally try to disentangle each channel's importance.

To estimate the magnitude of local general equilibrium effects, we depart from the difference-in-differences approach where we used non-oil households as our control group. Instead, we now study income and expenditure among exactly these households, and estimate how the share of oil-sector workers in a municipality affects the responses of non-oil households after oil price changes. There is systematic variation. In the top 95 percentile of municipal oil exposure, meaning that 7 percent of the working age population works in the oil sector, income of non-oil households responds by an additional 1.5 percent as compared to non-oil income in municipalities at the fifth exposure percentile where only 0.03 percent of working age population earn income from the oil sector. As compared to the income responses of the households with direct oil exposure, as displayed in Figure 1, the estimated general equilibrium effects are more delayed as they peak in years two and three after the initial oil price hike.

To illuminate whether it is the direct labor demand or Keynesian aggregate demand channels that underlie the GE estimates, we zoom in on different industries within the non-oil sector. In particular, we look at traded versus non-traded industries, motivated by the presupposition that Keynesian demand effects should have more influence on income in the latter group. No such pattern emerges. Moreover, we look at occupations within industries. Here we find that our baseline GE estimates primarily are driven by households in occupations that are more intensively employed in the oil sector. Overall, the GE effects seem to be primarily be driven by locally elevated demand for labor from the oil sector.

Notably, we also extend our analysis by replacing oil price changes with oil price shocks that are structurally identified. We here use the series for "speculative demand" shocks to the oil price by Kilian and Murphy (2014). Our main findings are not materially affected by using this shock series instead of the simple oil price changes.

Related literature Our paper is related to the small open economy macro literature which goes back to for instance Fleming (1962), Mundell (1963), and Dornbusch (1976). In this early literature, Keynesian demand multipliers play a prominent role in propagating external shocks and shaping how domestic policy should respond to them. Current textbooks still emphasize these channels, see for instance Dornbusch and Fischer (2017). Research has instead drifted toward micro-founded representative-agent models, as in Obstfeld and Rogoff (1995) or Galí and Monacelli (2005), which today arguably serve as the standard point of departure for studies of small open economy macroeconomics. These models leave little room for income hikes to spur consumption and thus set Keynesian style multipliers in motion. Auclert et al. (2021) elaborate on this point as they enrich the benchmark representative agent models with household heterogeneity and incomplete risk sharing, and show that this "sizes up the real income channel" of exchange rate movements. A key mechanism is that many households might be liquidity constrained and therefore consume a large fraction of the transitory income movements that exchange rate movements induce. Our results are indeed supportive this modeling approach, as we do find that households with direct labor market exposure to terms of trade shocks experience higher income and raise their expenditure considerably after a hike in export prices. Moreover, we do estimate spill-over effects to households who are not employed in the sector experiencing the price hike. But on the other hand, our results also suggest that these spillovers are primarily due to the direct "neoclassical" labor demand channel while the Keynesian aggregate demand channel is less important for propagation of export price changes. For the objective of modeling how external shocks transmit to small open economies, this pattern points toward frameworks that emphasize transmission via the labor demand of exporting firms.

Beyond the literature on small open economies, our expenditure analysis relates to the many recent studies that highlight how household heterogeneity matters for the transmission of macroeconomic shocks and policies in general, for instance Kaplan et al. (2018), Bayer, Born, and Luetticke (2020), and Auclert et al. (2020b). Part of this research agenda's appeal lies in its ability to connect structural models with empirical estimates of how households actually behave. This evidence typically comes from studies of how households respond to various sources of exogenous income variation, as in for instance Agarwal and Qian (2014), Di Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru, and Yao (2017), and Fagereng, Holm, and Natvik (2021). However, in theory it matters not just how households respond to past and current income shocks, but also how they respond in advance to anticipated income movements, as highlighted by Auclert et al. (2018). Evidence on such anticipation effects has so far been scarce, and here we make a contribution. To the best of our knowledge, the current work by Druedahl, Jensen, and Leth-Pedersen (2018) is the only other attempt to address this issue, as they study how households adjust consumption when they receive notice of future mortgage interest costs. They find that households who are not liquidity constrained respond when news arrives, while constrained households wait. We also find considerable anticipation effects, but in contrast to Druedahl et al. (2018) we do not find that this varies with households' initial liquidity. Our results are in this respect more in line with a standard permanent income hypothesis, but the magnitudes of expenditure responses that we estimate are larger than what a standard model of non-durable consumption would predict. We therefore speculate that durable consumption expenditure is key to explain our results, where planned durable purchases are shifted forward in time.[6]

An extensive literature uses time-series econometrics to structurally identify oil price shocks and estimate their effects on economic aggregates (Hamilton (1983), Kilian (2009), Baumeister and Hamilton (2019)). Overall, the evidence from this research suggests significant negative effects of oil price hikes on economic activity, and that it matters whether the price innovations are due to supply or demand side factors. However, to oil exporting economies oil price hikes have positive effects, as estimated by Bjørnland and Thorsrud (2016), Bjørnland et al. (2019) and Bergholt et al. (2019) for Norway. Our study is com-

[6]Our front-loaded consumption expenditure seems particularly at odds with the sticky expectations hypothesis set forth in Auclert et al. (2020b), where households are slow to realize the impact of macro variables and therefore do not increase consumption before actual income hikes materialize. Instead, from our estimates it seems that oil sector households anticipate that current oil price hikes will induce higher future income.

plementary to this literature, but differs as we attempt to assess some of the micro-level transmission mechanisms prominent in theoretical models.

Our focus on sectoral oil exposure is not novel. Keane and Prasad (1996) examines labor market outcomes across sectors, and finds that oil price fluctuations lead to changes in employment shares and relative wages across industries. Resting on regional variation in oil activity across US states, Acemoglu et al. (2013) estimate effects of regional income on health spending while Charles et al. (2018) estimate effects of local labor market tightness on disability take-up. The study of Kline (2008) is more closely related to ours, as he estimates how oil prices affect employment and wages in the US petroleum industry. The responses he finds in the US are remarkably similar to what we see in Norway. The wage response in the oil industry is hump-shaped: zero on impact, peak after 4-8 quarters, thereafter dampening. Kline (2008) also shows that there is significant inflow of workers into the oil sector after prices hike. However, he does not study effects on household balance sheets, expenditure or spousal income, nor attempt to quantify how households outside the oil sector experience general equilibrium effects like we do.

Our estimates of spousal income responses relate to the perennial literature studying family labor supply as insurance mechanism (Attanasio, Low, and Sánchez-Marcos, 2005; Blundell et al., 2016).[7] Evaluating the importance of several insurance mechanisms available to households, Blundell et al. (2016) find that spousal labor supply, both on the intensive and extensive margin, is a key insurance channel against wage shocks faced by the main income earner. Attanasio et al. (2005) show that female participation rates increase as uncertainty increases, and crucially that this effect on participation is greatest when the ability to borrow is limited, a point also noted by Lundberg (1985). Our results suggest that spousal labor supply is a significant factor in mitigating households' exposure to idiosyncratic income risk. Furthermore, we provide new evidence on the timing of the various insurance mechanisms: in the short term – under the assumption that the oil earners' wage increase

[7]Our attention to spousal earnings responses also connects with the large literature in labor economics focusing on the "added worker effect" (e.g. Lundberg (1985); Stephens (2002); Mankart and Oikonomou (2016)). The term added worker effect is, however, often related to the spouse's labor supply response to the job loss of the main earner within the household (Lundberg (1985)). Mankart and Oikonomou (2016) document that the added worker effect has been increasing since the mid-1990s. A recent contribution by Halla, Schmieder, and Weber (2020) echo previous studies on the added worker effect, with a small, but statistically significant effect. However, these spousal earning gains recover only a tiny fraction of the household income loss, and in the short term public transfers and taxes are a more important form of insurance. This contrasts with our findings, which suggests that spousal income responses provide substantial income insurance to the household. However, our analysis is clearly distinct from the traditional literature on added worker effects, for three reasons: (1) we do not condition our estimates on change in the employment status of the main earner (the oil worker), (2) we pool negative and positive shocks, and finally (3) we pool together the extensive and intensive margin responses of the spouse. We do this because our objective is not to estimate the exact margin of spousal earnings adjustment, be it wage, the intensive or the extensive margin of labor supply, but rather to estimate the quantitative importance of spousal income insurance for households.

is predictable – debt is the most significant instrument. In the longer term, we show that almost three quarters of the increased wage impulse of the oil workers is absorbed by the spouse. In a similar vein, Pruitt and Turner (2020) show that household income of two-person households displays significantly less procyclical skewness than individual income. Thus, Pruitt and Turner (2020) find that family labor supply mitigates cyclical income risk. Our findings also highlight the spousal response over time, as highlighted by Stephens (2002). Our findings show that on impact, spousal earnings drop more than the oil worker's income rises, suggesting a similar anticipation effect as in our expenditure estimates.

The spousal insurance margin plays a key role in the growing literature that evaluates the aggregate impact of family labor supply dynamics (Mankart and Oikonomou, 2016; Olsson et al., 2019; Bardóczy, 2020)). Particularly relevant to our study is the work of Bardóczy (2020), who develops a model where spousal income insurance acts as a powerful automatic stabilizer that significantly reduces consumption volatility over the business cycle. Our empirical evidence supports that the spousal income insurance margin is quantitatively important, and hopefully our estimates can help discipline future models with this channel.

The rest of this paper proceeds as follows. Section 2 presents key aspects of the Norwegian institutional setting, our data, and our empirical strategy. Section 3 presents our main results. In Section 4 we extend our analysis to assess robustness using identified oil price shocks instead of simple oil price changes and to illuminate the mechanisms underlying our main estimates, in particular the local general equilibrium effects. We then give a brief theoretical interpretation of our results in Section 5, before Section 6 concludes.

2 Institutional setting, data, and empirical strategy

In this section we briefly explain the role of oil exports in the Norwegian economy, describe our data, and present our empirical strategy.

2.1 Oil in Norway

Norway is an oil-exporting country. Although a relatively small supplier in the global scale, the Norwegian economy is highly oriented around oil- and gas production.[8] In 2014, oil and gas products amounted to 46 percent of Norways's total exports,[9] and according to official statistics from Statistics Norway direct employment in the petroleum sector amounted to three percent of total employment ($\approx$ 84 000 jobs as of 2014) (SSB 2017). In this paper, we will not distinguish between oil and gas activity, but instead refer to them jointly as "oil".

The Norwegian oil reserves are located off shore. Most of the production activity takes place along the western coast of Norway, but almost all municipalities (98 percent) have some employment in the oil sector. In addition, according to the Norwegian employers' association for oil and supplier companies in Norway, "Norsk Olje og Gass", within industries that more indirectly supply goods and services to oil production there are over 200 000 jobs.[10]

Profits from oil extraction are subject to resource rent taxation. In 2014, approximately 27 percent of all the Norwegian state revenues came from oil.[11] These revenues are invested in a sovereign wealth fund, the Government Pension Fund Global, commonly known as the oil fund. The fund invests in foreign assets only. Since 2001, the government's saving and spending decisions have been separated through a budgetary rule, stating that the budget deficit excluding oil-tax income cannot exceed a given fraction of the fund's current market value. This fraction was initially 4 percent, but has later been cut to 3 percent.[12]

The Norwegian oil exposure is also reflected in the exchange rate, which tends to move together with the oil price.[13] In short, the Norwegian economy relies heavily on oil exports, with the implication that Norwegian household income is likely to be exposed to fluctuations in the global price of oil.

[8] Norwegian oil and gas supplies constitute about two and three percent of total global oil and gas supply, respectively.

[9] https://www.norskpetroleum.no/produksjon-og-eksport/eksport-av-olje-og-gass/

[10] Norsk Olje og Gass (2021, see https://www.norskoljeoggass.no/en/).

[11] https://www.norskpetroleum.no/okonomi/statens-inntekter/

[12] The Norwegian budgetary rule is titled "Handlingsregelen". Further details on the public strategy for handling oil revenues are provided by Bjørnland (2018).

[13] The correlation between the change in the average oil price and the change in the average NOK/USD exchange rate in our sample period was -0.63. In-depth analysis of the link between oil price and NOK fluctuations is provided by for instance Akram (2004).

2.2 Data

Our empirical analysis uses administrative data from Norway, covering the universe of tax-paying residents. Several different datasources are linked at the individual level, and are thereafter de-identified before we access them. Most of our data are maintained by Statistics Norway. Our earliest data start in 1993 and ends in 2015. After merging the various data needed for our purposes, the final dataset run from 1996 to 2014.

We utilize two main data sources. Our first main source is the balance sheet of individuals, collected by the tax authority ("Skatteetaten") for tax purposes. This includes earned income (salaries and business income), capital income (interest income and expenditure, realized capital gains, dividends, government transfers and taxes), as well as all wealth components (housing wealth, bank deposits, stocks, bonds and mutual funds, and total debt).[14] This information includes links between married and cohabitant individuals, enabling us to aggregate data to the household level.

Based on these data we construct an imputed measure of household-level yearly consumption expenditure. We build on the procedure of Fagereng and Halvorsen (2017), later improved and used by Fagereng et al. (2021) and Holm, Paul, and Tischbirek (2021). The approach rests on the budget identity that consumption expenditure equals income minus active saving, where the latter is the part of income that is set aside for wealth accumulation during the year. The main challenge with this approach is to parse out the contribution of unrealized capital gains, i.e. asset price changes, to wealth accumulation over the year, sometimes referred to as "passive saving". Our approach essentially assumes that households rebalance their portfolio at the end of the year only, allowing us to use observed anual price changes to subtract the contribution from price growth on the beginning-of-year asset holdings from the annual wealth change. In practice, this procedure is important for the wealthiest households only, since the vast majority of Norwegians hold a small portfolio share of other financial assets and liabilities than interest-bearing deposits and debt for which we directly observe yields (interest income and expenditure are directly observed). For further details on our measure of consumption expenditure, see the references above.[15]

Importantly the imputation procedure outlined above does not distinguish non-durable from durable consumption expenditure. We therefore refer to the resultant measure as "expenditure" hereafter.

In addition to balance sheet information, we have rich data on individual characteristics, including education, sex, age, number of children and municipality of residence. In most

[14] Because all Norway levies wealth taxes, end-of-year holdings in all asset classes are reported.

[15] Other studies that impute consumption expenditure from registry data include Browning and Leth-Petersen (2003), Eika, Mogstad, and Vestad (2020), and Kolsrud, Landais, and Spinnewijn (2020). A recent evaluation of the approach is provided by Baker, Kueng, Meyer, and Pagel (2021).

of our empirical analysis our unit of observation is the household, although we also use individual level variables to dissect the movements in household income. We focus on working age individuals (i.e. individuals between 25 and 67 years old). In addition, we make some minor sample restrictions to make sure that our results are not driven by outliers. These are as follows. We exclude a household entirely from our sample (i.e. all years) if they in any year during our sample period (1) earn private business income, or (2) the value of its stocks, bonds and mutual funds is in the top five percent. Second, we exclude a household-year observation if (3) the imputed expenditure in that year is below zero or in the top 5%, or; (4) disposable income is less than one social security base unit in that year ($\approx$ \$11000 in 2014).[16]

Our second main data source is linked employer-employee data at the individual level. These data contain information on individuals' employee-employer relationship, including duration of employment, expected working hours per week, pay, the municipality in which the firm is located and a 5-digit NACE code. The latter is an international standard of industrial classification and categorizes firms by their production. This sector classification allow us to identify which sector each individual is working in, in any given year. From 2005, the data also classifies the individuals' occupation from a list of roughly 90 categories (e.g. engineer, teacher, office clerk).

We use the employer-employee data to characterize individuals' labor market exposure to oil price fluctuations. Here we implement Statistics Norway's categorization of the oil industry in Norway, based on employers' NACE codes (Statistics Norway, 2017). NACE codes are divided into "Oil Extraction and Related Services" (henceforth "oil extraction", or "sector 2") and "Construction/Furnishing of platforms and Supply Bases" (henceforth "oil service", or "sector 1"). The remaining NACE codes are categorized as "non-oil" (or "sector 0"). We use this classification to assign each individual to a sector, equal to the sector indicated by their employer's NACE code. We consider the three industries' exposure to oil price fluctuations in an ordinal manner:

$$Sector\ 2 > Sector\ 1 > Sector\ 0$$

Statistics Norway's definition of the oil service sector is strict, as it only contains industries where output is more or less a direct input to production in the oil extraction sector. Thus, some firms that in reality produce for the oil sector, but where their NACE-defined sector in total only produces a small fraction of its output for the oil sector, are not classified as belonging to the oil service sectors.[17] Our classification therefore does not capture

[16]Norwegian social security entitlements are typically denominated in units of the basic amount "G", which in turn is yearly adjusted to track average income.

[17]Statistics Norway mention e.g. office supplies, cleaning services etc. as sectors where individual firms

firms that indirectly produce for the oil industry, or where only part of their production is related to oil. The benefit of our measure is that it has been subject to comprehensive quality control in Statistics Norway. The drawback is that we are likely to classify some individuals as non-oil workers, when in fact their income is somewhat directly exposed to oil price changes.[18]

Once individuals are categorized, we classify households into the same three categories. We classify households as follows:

$$Sector_{ij}^{h} = max\{Sector_i, Sector_j\}$$

where i and j are individuals in household h. That is, household h is classified according to the most oil exposed household member i or j.

Figure 2 shows the sector shares relative to total household observations each year. It shows that oil sector households (red bars) amounted to approximately one percent of all households in our first sample year, 1996. Over the next eighteen years, this share grew up to roughly 2.7 percent, reflecting the increasing importance of the petroleum sector in the Norwegian economy during these years. The oil service sector (blue bars) followed the development in the oil sector closely up to 2008. However, in 2009, Statistics Norway reclassified a large share of businesses from the oil service to the oil sector, creating a clear break in the two series plotted in Figure 2.[19] All-in-all the two sectors together amounted to 3.25 percent in 2014, up from 2.0 and 2.75 percent in 1996 and 2008, respectively.

realistically might deliver to the oil industry, but where the sector as a total does not, and are thus not classified as "oil service".

[18]In Section 2.3 we discuss the potential concerns this creates for our identification strategy, and how we deal with it.

[19]A robustness analysis of the period up to 2008 indicates that this shift is not important for our results.

Figure 2: Share of households classified as oil households or oil service households

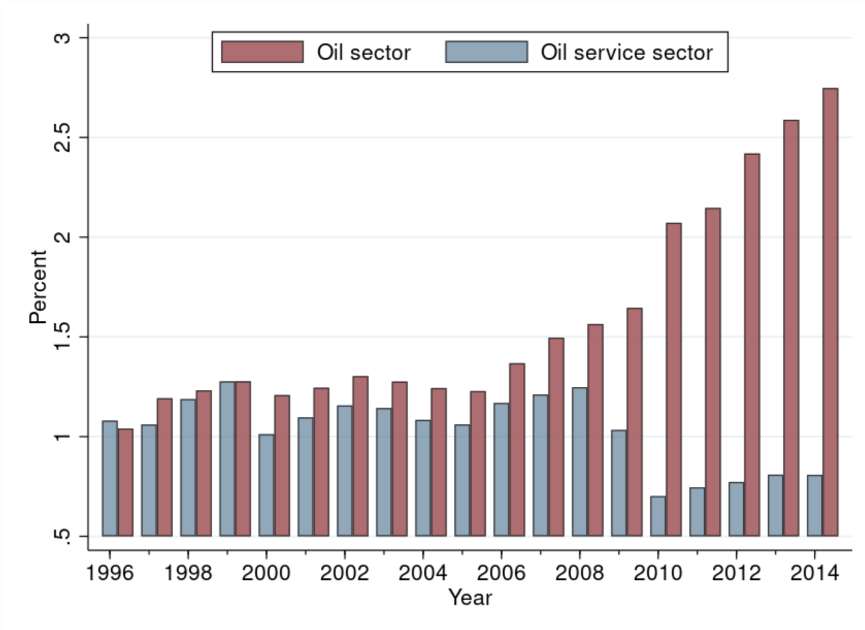

Notes: The figure reports the percent share of households classified as oil households (red bars) or oil service households (blue bars) in our main estimation sample. The classification of sectors follows Statistics Norway (2017), see Section 2 for details.

Our second measure of oil exposure is based on municipality of residence, and the share of that municipality's individuals who work in either of the two oil sectors. Our simple measure of local exposure' to oil price fluctuations is the share of oil workers in the municipality. Households are then assigned values equal to their municipality's share of oil workers each year (t):

$$OilShare^k_{it} = OilShare^k_t = \frac{N^k_t\{Sector = 1\} + N^k_t\{Sector = 2\}}{N^k_t}$$

where k is municipality, i is household.

Our empirical strategy rests on the variation in the local exposure variable. Figure 3 (a) plots the density of all municipality-years on the regional exposure variable (with N = 8208). It shows that most municipalities cluster in the lower part of the distribution, with roughly 80 percent of municipality-years having less than two percent oil workers. Next, the figure shows a long right tail, with a maximum of 27 percent oil share. In Figure 3 (b) we graph the top ten percent municipality-years with respect to oil-worker share (N = 838). In this plot, we count the frequency at each share. This plot shows a fair amount of observations between four and fifteen percent.

Figure 3: The distribution of local exposure

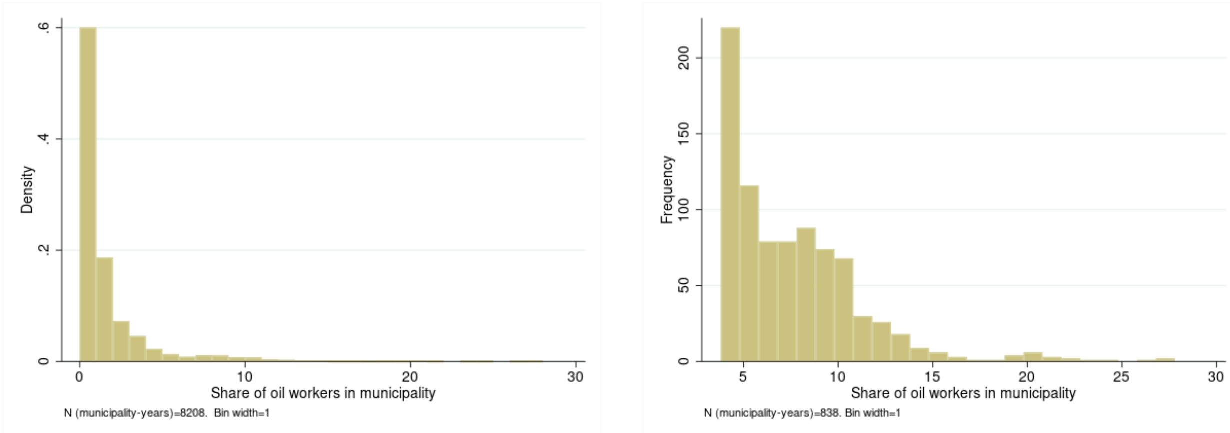

(a) All municipality-years

(b) Top-10% municipality-years

Notes: Panel (a) shows the distribution of oil-worker shares across all municipality-years in our sample. Panel (b) reports the frequency of municipality-years within the top 10 percent of the oil-worker share distribution (shares above ≈3.7 percent). The share of oil workers is calculated using individuals as the unit of observation. The classification of sectors follows Statistics Norway (2017), see Section 2 for details.

Summary statistics for our sample are displayed in Table 1. The table shows that there is a level difference between oil, oil service and non-oil households with respect to income, expenditure and wealth. On average, oil service and oil households have higher earned income compared with non-oil households. Disposable income is also higher, but less so due to progressive taxes. Oil service and oil households are also more wealthy compared with non-oil households, where housing wealth is the most important component. Oil dependent households also hold more debt. Imputed expenditure is also considerably higher for the oil dependent households. Note that our imputed expenditure includes an estimate of housing consumption, which leads to a higher average expenditure than disposable income.

Table 1: Summary statistics

	Sector					
	Non-oil		Oil Service		Oil	
	Mean	s.d	Mean	s.d.	Mean	s.d.
Earned income	70 588	(57 639)	114 995	(56 366)	168 872	(83 031)
Disposable income	64 777	(35 926)	88 366	(38 144)	120 296	(53 994)
Expenditure	74 332	(131 605)	96 841	(86 191)	128 460	(99 649)
Housing wealth	437 937	(508 253)	526 461	(480 095)	707 956	(610 918)
Deposits	24 252	(51 013)	28 505	(53 108)	42 656	(74 568)
Other assets	2 671	(7 522)	3 983	(9 062)	7 922	(12 852)
Debt	121 964	(159 652)	175 445	(175 282)	263 215	(230 143)
Net wealth	342 896	(481 575)	383 503	(452 234)	495 319	(570 792)
Observations	18 048 885		191 373		306 733	
Percent of total	97.31%		1.03%		1.65 %	

Notes: The table summarizes the mean and standard deviation of households' balance sheet components, by sector of employment, within our main estimation sample. The sample period is from 1996 to 2014. Individuals included in our sample are between 25 and 67 years old. Disposable income is defined as the sum of labor income, transfers and capital income, after tax. Expenditure is imputed from the household budget constraint (see e.g. Fagereng and Halvorsen (2017)).The classification of sectors follows Statistics Norway (2017), see Section 2 for details. We report monetary variables in USD, 2011 prices. Conversion from NOK to USD applies the average exchange rate between USD and NOK in 2011, NOK/USD = 5.77.

Our empirical analysis relies on differences between the groups with respect to oil exposure. In our main analysis we use all non-oil households as the reference group. For the purpose of our research question, namely how oil price fluctuations transmit to households, we should not exclude specific groups from the reference category. However, as an extension to our main analysis we zoom in on how oil households respond relative to various subgroups of non-oil households. A priori, we expect that when we compare oil workers to control groups within similar occupations outside the oil sector, the differential income responses will be mitigated and maybe even disappear entirely (and vice versa for more dissimilar groups), as these non-oil households compete in the same labor market as the oil workers. To investigate this hypothesis, we use data on occupational codes to define a subset of non-oil households that work in occupations that are frequently observed in the oil industry. Section 4.3 presents a more detailed description of the data and the results.

2.3 Empirical strategy

Our empirical analysis uses the heterogeneity in workers' exposure to the oil market described in the previous section, and that fluctuations in the price of oil are determined in the global market. We will estimate responses to oil price changes at the individual and household level, given pre-determined exposure to oil. To this end, we apply local projections as proposed by Jordà (2005), with a shift-share type design as in Chodorow-Reich et al. (2021).

Our starting point is the following specification:

$$\begin{aligned}\frac{Y_{it+h} - Y_{it-1}}{Income_{it-1}} =& \alpha_i + \sum_{j=1}^{2} \boldsymbol{\beta}_j \mathbf{X}_{it-j} + \gamma_1 \Delta Oilprice\%_t \\ &+ \gamma_2 EXP_{it-1} + \delta^{EXP}_{t+h} \Delta Oilprice\%_t \times EXP_{it-1} + e_{it}, \end{aligned} \tag{1}$$

which also is the regression model that lies behind Figure 1 in the Introduction. $\frac{Y_{it+h}-Y_{it-1}}{Income_{it-1}}$ is the cumulative change in Y from year $t-1$ up to the end of year $t+h$ for household i, $\Delta Oilprice\%_t$ is the percentage change in the oil price from the last day in $t-1$ to the last day in t, and EXP_{it-1} is entity i's predetermined degree of exposure to the oil price change. In our first set of results this categorical variable takes on three values (oil-, oil sevice-, and non-oil households) describing household i's sector of employment, as explained in Section 2. EXP_{it-1} is then a vector of indicator variables categorizing households by their sector of occupation. In contrast, when estimating spillover effects (household i's indirect exposure to oil), EXP_{it-1} is continuous, and measures the share of oil workers in household i's municipality of residence (again, see Section 2). $\Delta Oilprice\%_t \times EXP_{it-1}$ is our "shift-share variable". It is the interaction term between the oil price change and household i's pre-determined oil exposure. Lastly, our controls in Equation 1 include $\mathbf{X}_{it-j}$, which is a vector of pre-determined time-varying control variables, and individual-fixed effects, α_i.[20] Note that our left-hand side variable is always normalized with $Income_{it-1}$. By scaling with the same factor, we make our coefficient estimates easily comparable across the different outcome variables, Y.

The coefficients of interest are in the vector δ^{EXP}_{t+h}. They reflect how much households belonging to the oil or oil-service sector group are affected by the oil price change, relative to the reference group, which will be the non-oil households. The coefficient γ_2 picks up average differences in outcomes across sectors, such as those that we saw in Table 1 and would be subsumed in the individual-fixed effects α_i if households did not switch sectors. Since households actually do change sectors over our 20-year sample period, these effects

[20] As is standard in these types of local projections, the vector $\mathbf{X}_{it-j}$ contains lags of our shift-share variable, the oil price change, and the one-year growth of Y at lag j. $\boldsymbol{\beta}_j$ is the associated vector of coefficients. In our main specifications, we use $j = 2$.

must be controlled for separately.

The estimates of δ_{t+h}^{EXP} resulting from running the regression model (1) were plotted in Figure 1 in the Introduction. In that figure we saw that after an oil price hike, household income in Norway tends to rise not only in the directly oil-related sectors, but in other sectors too. This could be due to indirect effects from the oil price change spilling over across sectors. But equally likely, it reflects that the oil price is endogenous to other global macro economic factors. For instance, U.S. GDP growth tends to affect and be affected by the oil price (Baumeister and Hamilton, 2019; Kilian, 2009). We thus cannot say if the estimates from a specification like (1), such as those displayed in Figure 1, are due to oil prices or some other macro phenomenon occurring simultaneously.

Moreover, there could be general equilibrium effects operating at the local level. Oil-sector workers tend to live in the same municipalities, as seen in Figure 3. Hence, if there are spillovers from oil sector workers to the local economy, then the magnitude of general equilibrium effects will be stronger in exactly those regions where oil-sector workers live. Therefore, even if we consider the difference between oil and non-oil workers, we cannot immediately say how much of such a difference is due to oil price fluctuations, or local general equilibrium effects. An example of such an effect could be that local house prices respond more in regions where there are more households with income from oil-related activities. In the third part of our estimates, we are interested in quantifying exactly such general equilibrium effects, but in the first sets of estimates we want to filter them out so as to zoom in on what plausibly constitutes household-level responses to household-level income.

Global endogenous factors and local business cycles therefore make it difficult to give the results from regression model (1) a causal interpretation. To control for such confounding factors, we include municipality-time fixed effects in our specifications and focus on the differences between households with different exposure to oil prices.[21] After including municipality-time fixed effects, our baseline specification then takes the form

$$\begin{aligned}\frac{Y_{it+h} - Y_{it-1}}{Income_{it-1}} =& \alpha_i + \tau_{m,t} + \sum_{j=1}^{2} \boldsymbol{\beta}_j \mathbf{X}_{it-j} + \gamma_2^{EXP} EXP_{it-1} \\ &+ \delta_{t+h}^{EXP} (\Delta Oilprice\%_t \times EXP_{it-1}) + e_{it} \end{aligned} \tag{2}$$

where $\tau_{m,t}$ is the municipality-time fixed effect for municipality m in year t, and the oil price change no longer enters separately as it is subsumed by the time-fixed effects. Since our regressions include municipality-time fixed effects, the average effects across households

[21] Another approach is to use identified "oil price shocks", i.e. oil price fluctuations that are (arguably) purged of such endogenous factors. We do this in Section 4. There we use speculation-induced oil price shocks as identified in a structural VAR model (collected from Baumeister and Kilian (2016a)) instead of the much simpler percent change in the oil price.

are canceled out, and we are left with interpreting the differences between oil and non-oil households.

To give our estimates a causal interpretation we rest on two identifying assumptions. First, oil price changes need to be exogenous, and not predicable. Since oil prices are determined in the global market where Norway is a minor player, and follow approximately a random walk, we argue that this assumption is reasonable. Our pre-treatment responses (see Figure 16 - 17 in the Appendix), support this assumption. One cannot see systematic pre-treatment differences ("pre-trends") between oil and non-oil households that correlate with the oil prices changes. Second, variables that are correlated with oil price changes, after controlling for $\mathbf{X}_{it-j}$, α_i and $\tau_{m,t}$), cannot affect households in the different oil-exposure groups in a systematic manner. For example, if government budgets rise with rising oil prices, and larger government budgets tend to favor or disfavor oil households relative to non-oil households, this would constitute a breach to the second assumption. Our extension in Section 4, using structurally identified oil price shocks instead of oil price changes, lend support to our claim that both identifying assumptions hold. Our baseline results and these results are largely overlapping.

Beyond the nature of our shift-share variable, our strict classification of sectors is a potential threat to our empirical strategy. Our classification is likely characterizing some households as "non-oil" while they in reality earn some of their income directly from or in service to the oil industry. The reason is that firm NACE codes are based on firm production, but not on the individual firms' clients, and in Statistics Norway's (and our) classification, only NACE codes where output is oil or enters as factors of oil production, are classified as oil or oil service sectors. However, clients may vary within the same NACE code, so that firms in the same code can deliver services or products to the oil industry, and others not. This misclassification error might be important when we estimate local spillover effects. If our misclassification error is correlated with the oil intensity of municipalities, this will create a positive bias in our spill-over estimates because it will include the positive (direct) effect on the oil households who are misclassified as non-oil households. For our other analyses, where we compare oil to non-oil households, the same error will bias our estimates *downward*. In Section 4.4 we explore the severity of this issue, by focusing on households employed in industries where such misclassification is unlikely, such as retail. We find similar spill-over effects as in our baseline estimates, and therefore conclude that the bias is likely to be small.

3 Main Results

We split the exposition of our main results into three. First, we concentrate on the profile of spending responses relative to income responses after oil price changes, with the ambition to illuminate intertemporal propensities to consume out of income changes that are predictable, but have not yet materialized. To this end we focus on household income, as our expenditure measure only applies at the household level. In the second part we instead distinguish between the earnings responses of each party in a couple, with the objective to estimate individuals' propensity to adjust their earnings in response to income variation for their spouse. Third, we attempt to quantify the importance of general equilibrium effects for non-oil workers after an oil price hike.

3.1 Household income responses and the intertemporal propensity to spend

Figure 4 displays how income and expenditure respond among households with earnings from the oil service sector, panels (a) and (c), or from the oil sector, panels (b) and (d), using specification (2) which controls for municipality-time fixed effects. The same plots with pre-trends included are displayed in Figure 16 and 17. Period 0 is the year in which the oil price change occurs.

We see that both oil-sector and oil-service-sector incomes increase in a hump-shaped manner, with negligible movement in the impact year. In contrast, expenditure responds immediately, as seen in the lower panels.[22]

If we assume these income responses are predictable for the households, the expenditure responses reflect anticipation effects and can be used to estimate intertemporal propensities to spend out of anticipated future income changes. As oil price changes are salient to the public, and because the income responses we estimate are the systematic average responses over a 20-year period, we believe such an interpretation is warranted. Moreover, the delayed response of income is a likely consequence of the staggered nature of Norwegian wage setting, where wages are typically set in the late Spring each year. Hence, it should not come as a surprise to the workers when their income responds to oil price changes that materialized over the previous year.

[22]The results reported are for regressions with an unbalanced panel. We report the results with a balanced panel in the Appendix, Figures 20 and 21. Results are very similar.

Figure 4: Income and expenditure responses to a 10 % change in the oil price. Oil-sector households and oil service sector relative to non-oil households.

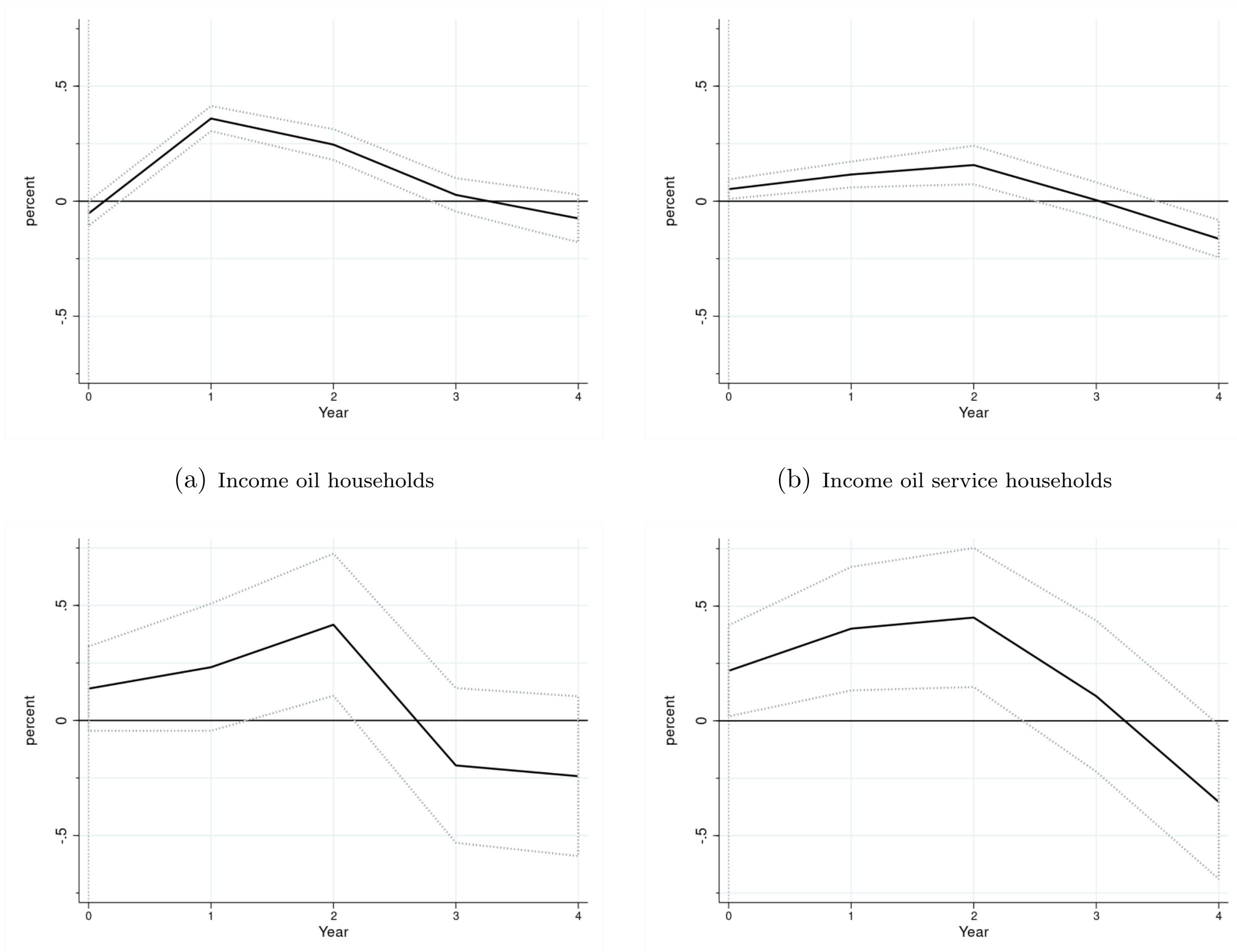

(a) Income oil households

(b) Income oil service households

(c) Expenditure oil households

(d) Expenditure oil service households

Notes: The figure plots the responses of household income and expenditure in every period $t+h$ relative to income in period $t-1$, after a 10 percent increase in the oil price in period t. Dotted gray lines are the 95 percent confidence bands. The plots follow from the estimated coefficients δ_{t+h}^{EXP} in Equation 2, with $EXP_{t-1} = sector_{t-1}$ as exposure variable. Disposable income is defined as the sum of labor income, transfers and capital income, after tax. Expenditure is imputed from the household budget constraint (see e.g. Fagereng and Halvorsen (2017)). All variables are scaled by disposable income in $t-1$. The classification of sectors is based on Statistics Norway (2017), see Section 2 for details.

We next explore how the anticipation effect in spending is financed in Figure 5. We here focus on oil-sector households only, while similar responses for oil-service sector workers are reported in the appendix, Figure 17. The main insight is that the additional expenditure in the first year is primarily financed with debt and only to a negligible extent by drawing down deposits. Thereafter, as income materializes, debt is repaid. Note that our controlling for municipality-time fixed effects means these movements are not driven local general equilibrium effects. Figure 18 in the appendix illustrates the role of our municipality-level controls effects, by displaying impulse responses with time-fixed effects only.

Figure 5: Debt and deposit responses to a 10 % change in the oil price. Oil households relative to non-oil households.

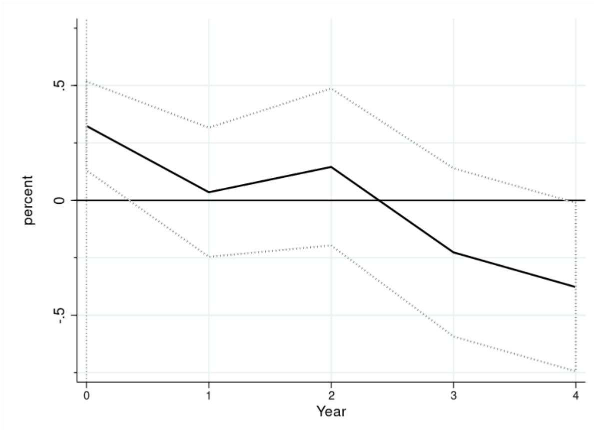

(a) Debt

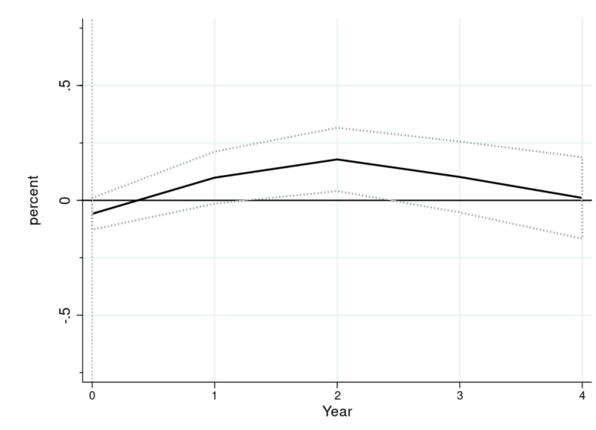

(b) Deposits

Notes: The figure plots the responses of household balance sheet components (as reported in each sub-figure caption) in every period $t+h$ relative to income in period $t-1$, after a 10 percent increase in the oil price in period t. Dotted gray lines are the 95 percent confidence bands. The plots follow from the estimated coefficients δ_{t+h}^{EXP} in Equation 2, with $EXP_{t-1} = sector_{t-1}$ as exposure variable. All variables are scaled by disposable income in $t-1$. The classification of sectors is based on Statistics Norway (2017), see Section 2 for details.

As mentioned in the introduction, households' intertemporal propensities to consume out of future income movements are key to how shocks propagate through the economy in macroeconomic models. This point is laid out in particular detail by Auclert et al. (2018). The dynamic responses of spending and income in Figure 4 are informative about these moments. However, the future income innovations we observe do not pertain to one period only, but are persistent and last several periods. We therefore cannot ascribe the spending response in, say, period 0 to the income response in some particular period $t+j$ only.[23] Instead, we can translate the impulse responses in Figure 4 into how expenditure cumulatively responds to cumulated income movements over time. We do this in Table 2. The top rows display the cumulative expenditure an income responses, which simply are the sum of the period-by-period responses in Figure 4. The estimated anticipation effect for expenditure in period 0 of 0.14 is approximately half of the income response in period

[23]This is also why we cannot use the oil price change in period t as an instrument for the income change in any specific period $t+j$. The fact that income responses take place over several periods means that the exclusion restriction would be violated with such an approach.

1, while in period 3, the two cumulative effects are almost identical. After 4 periods the expenditure drop we saw in Figure 4 means that the cumulated increase in spending is only 0.35, while it is 0.5 for income. That is, 4 years after the oil price change, expenditure has increased by 70 percent of the total income increase.

Table 2: The responses of spending and income over time

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
$\sum_{j=0}^{h} \Delta C_{t+j}$	0.14	0.37	0.79	0.59	0.35
$\sum_{j=0}^{h} \Delta Y_{t+j}$	-0.05	0.31	0.55	0.58	0.50
$\frac{\sum_{j=0}^{h} \Delta C_{t+j}}{\sum_{j=0}^{4} \Delta C_{t+j}}$	0.40	1.06	2.26	1.69	1.00
$\frac{\sum_{j=0}^{h} \Delta C_{t+j}}{\sum_{j=0}^{4} \Delta Y_{t+j}}$	0.28	0.74	1.58	1.18	0.70

Notes: The statistics in this table are based on the estimated coefficients δ_{t+h}^{EXP} in Equation 2 for the oil sector relative to the non-oil sector. The table reports the cumulative sum of estimates from $h = 0$ to $h = j$ for expenditure (row one) and disposable income (row two). Row three reports the cumulative expenditure response up to period j out of the total response up to $j = 4$. Row four reports the cumulative expenditure response up to period j out of the total income response up to $j = 4$.

The bottom two rows of the table serve to more precisely convey the timing of spending and income changes. The third row shows how much extra has been spent in period j relative to the total effect after 4 years. We see that 40 percent of the total response takes place immediately, and that after two more years the cumulative response is more than twice as high as it is after 4 periods. Finally, the fourth row reports the period j cumulative expenditure effects relative to the cumulative income response 4 years beyond the impact year. We see that the impact-year response of expenditure is 28 percent that of the cumulative 5-year income response.

These responses indicate not just that expenditure reacts immediately to anticipated income changes, but also that the expenditure response overshoots before contracting thereafter. What could explain such overshooting? As we return to later, durable consumption goods are one likely explanation. If a household experiences a positive shock to expected future income, it will likely want to consume more of both durable and non-durable goods. In theory, this would typically lead to a disproportionately high increase of durable expenditure in the short run. Moreover, planned future purchases of durables would tend to be shifted forward in time, which could lead to a negative expenditure response at longer horizons. Hence, our results point to a durable goods model, as we return to in Section 5.

3.2 Spousal income responses

In many households there are two adults of whom only one works in the oil or oil-service sector, whereas the other obtains his or her income elsewhere. This allows us to explore how oil price induced income changes for oil-sector workers transmit to the income of their spouses.[24] Hence, we can shed light on the degree to which couples insure against idiosyncratic income risk along the spousal earnings margin, in the sense that one spouse earns less when the other spouse earns more. Our intention here is not to estimate the exact margin of spousal earnings adjustment, be it wage, the intensive or the extensive margin of labor supply, but rather to estimate the quantitative importance of spousal income insurance for households.[25]

To this end, we use the same local projection framework as described in equation (2). We concentrate on households where there are two adult members, and drop the few couples where both work in the oil or oil-services sector. The categorical exposure variable EXP_{it-1} now identifies those households where only one member works in the oil sector in year $t-1$. We consider income at three levels: the household (as before), the oil-sector worker ("oil income"), and the spouse who is not in the oil sector ("non-oil income"). For each level, the individual-fixed effects in (2) are adjusted correspondingly, so that they are estimated at the household or the actual individual household-member level. Moreover, we separately estimate responses for gross income including government transfers ("total income") and for earnings only.

As before, our interest lies in the estimates of δ_{t+h}^{EXP} and we report them in Figure 6. Panel (a) reports the responses of total income within the household, and panel (b) reports estimates for earnings only.

[24] Importantly, in Norway income taxes are determined at the individual and not at the household level. This is important for our purposes, as it means there is no effect of oil-sector workers' income on their spouses' income tax.

[25] With a unitary model of the household in mind, this exercise will pick up the effect of household-level income on labor supply, or the "marginal propensity to earn" out of an income hike in the terminology of Auclert, Bardoczy, and Rognlie (2020a).

Figure 6: Income and earnings within the household after a 10% oil price change.

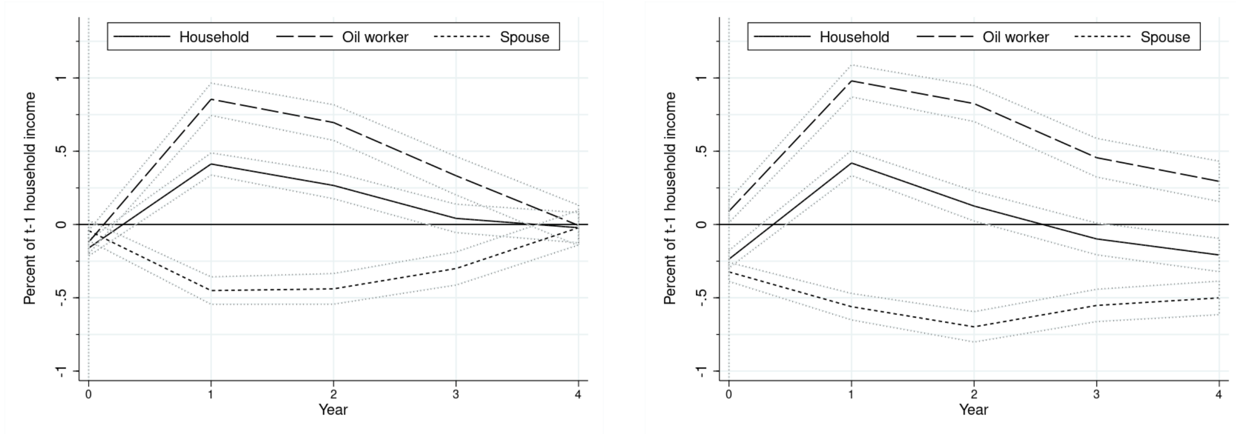

(a) Total income

(b) Earnings

Notes: Panel (a) plots the responses of total after-tax income from earnings and transfers at the level of the household (solid line), the oil-sector worker in the household (dashed line) and the spouse outside the oil sector (dotted line), after a 10 percent increase in the oil price in period t. Panel (b) plots the responses at the same three levels for earnings only. The plots follow from the estimated coefficients δ_{t+h}^{EXP} in a specification of the type in equation (2), where EXP categorizes whether one or none of the household members were emplyed in the oil sector in tear $t-1$. Dotted gray lines are the 95 percent confidence bands. All variables are scaled by disposable income in $t-1$. The classification of sectors is based on Statistics Norway (2017), see Section 2 for details.

The solid curve in panel (a) depicts the same income concept as seen earlier, but excluding taxes and net interest income. Its trajectory is essentially the same as before, and the minor differences are due to the fact that we are considering a slightly narrower sample than previously. The upper dashed curve is the income obtained by the oil-sector worker in the households. As we would expect, this response lies above that of household-level income. They both converge back to their previous levels 4 years after the price hike took place. The key outcome is the bottom dotted line, which shows the income response for the spouses of oil-sector workers. We see a marked negative income response for these individuals. Qualitatively, their income response is approximately symmetric to that of the oil-workers, only in the opposite direction. The movement is negligible in the impact year, strengthens gradually, and then reverts. 4 years after the impulse, spousal income has reverted.

Panel (b) shows that a similar pattern holds for earnings. Here, the response for the oil-workers is somewhat higher and that for the non-oil spouses somewhat more negative than in panel (a). The gap between total income and earnings responses is due to government transfers. The negative household-level earnings response after 4 years thus reflects that these households are getting slightly higher transfers than they did before the oil price innovation.

Quantitatively, Figure 6 (a) shows that for every dollar of extra income from the oil-workers, the non-oil spouses' income falls by more than one half. Table 3 reports our point estimates for the cumulative income responses inside the household, and the lower line in the table scales the non-oil spouse's income response by that of the oil-worker. In the bottom right corner, we see that 4 years after an oil price hike, the average non-oil spouse has cut his or her income by approximately 70 cents to the dollar of the oil-workers cumulative income hike.

Table 3: The spousal income response

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
$\sum_{j=0}^{h} \Delta Y_{t+j}^{spouse}$	-0.04	-0.49	-0.93	-1.23	-1.25
$\sum_{j=0}^{h} \Delta Y_{t+j}^{oil}$	-0.12	0.73	1.42	1.76	1.76
$\frac{\sum_{j=0}^{h} \Delta Y_{t+j}^{spouse}}{\sum_{j=0}^{4} \Delta Y_{t+j}^{oil}}$	-0.02	-0.28	-0.53	-0.70	-0.71

Notes: This table reports the cumulative income responses of spouses of oil workers (row 1), the cumulative income response of oil workers, and the cumulative income responses of spouses divided by the total income response of oil workers over five years. Income include individual earned income and transfers.

3.3 Regional exposure and spillovers to non-oil households

Oil activity in Norway is geographically concentrated, as explained in Section 2. Hence, households with income from oil activities tend to co-locate in certain municipalities, and we would expect that oil price increases induce stronger local general equilibrium effects in those areas where more oil-sector households reside.

In our baseline estimates above, we deliberately parsed out such geographical variation by controlling for municipality-time fixed effects in our regressions. Now, we instead switch attention to quantifying how strong the local general equilibrium effects might be. As an illustration, Figure 7 therefore reports income responses within the household with time-fixed effects that are common to all and not specific to each municipality. That is, we replace $\tau_{m,t}$ with τ_t in equation (2). The results do indeed suggest that the effects of oil price changes spill locally over to non-oil incomes. When we impose common time fixed-effects across all municipalities, the income response of households with one member working in the oil sector increases somewhat. This change is due to a moderation of the non-oil spouse's income decline, not the oil-sector worker. Quantitatively, the negative spousal earnings response is dampened by approximately one fifth when we let the local variation in general equlibrium effects factor in.

Figure 7: Income and earnings within the household after a 10% oil price change, the role of municipality-time fixed effects.

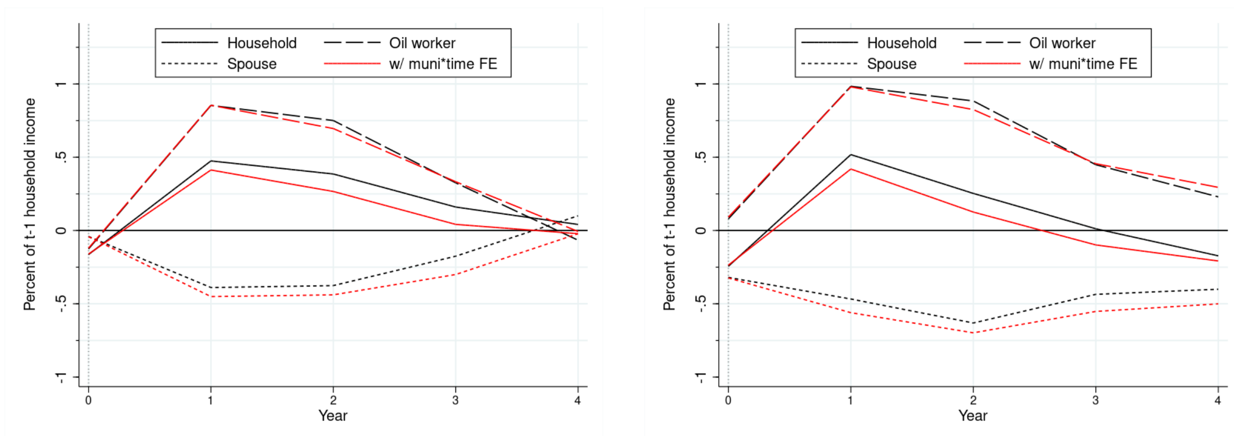

(a) Total income

(b) Earnings

Notes: Panel (a) plots the responses of total after-tax income from earnings and transfers at the level of the household (solid line), the oil-sector worker in the household (dashed line) and the spouse outside the oil sector (dotted line), after a 10 percent increase in the oil price in period t. Panel (b) plots the responses at the same three levels for earnings only. The black curves follow from specifications with time-fixed effects only, the red curves are the baseline estimates with municipality-time fixed effects. The plots follow from the estimated coefficients δ_{t+h}^{EXP} *in a specification of the type in equation (2), where* EXP *categorizes whether one or none of the household members were employed in the oil sector in tear* $t-1$*. All variables are scaled by disposable income in* $t-1$*. The classification of sectors is based on Statistics Norway (2017), see Section 2 for details.*

The same qualitative patterns as in Figure 7 emerge when we redo our baseline exercises from Section 3.1 with time- instead of municipality-time fixed effects, consistent with local spillover effects to households outside the oil sector.[26] However, these patterns are merely suggestive, hardly statistically significant, and do not quantify how strong the local general equilibrium effects might be. To assess the strength of these effects, we estimate how the responses of non-oil households vary with their municipality's exposure to the oil sector. We

[26]A comparison of estimates with time-fixed and municipality-time-fixed effects is given in the appendix, section A.2.

use the following econometric specification:

$$
\begin{aligned}
\frac{Y_{it+h} - Y_{it-1}}{Income_{it-1}} =& \alpha_i + \tau_t + \sum_{j=1}^{2} \boldsymbol{\beta}_j \mathbf{X}_{it-j} + \gamma OilShare_{imt-1} \\
&+ \delta_{t+h}(\Delta Oilprice\%_t \times OilShare_{imt-1}) + e_{it},
\end{aligned} \tag{3}
$$

where $OilShare_{imt-1}$ is the share of workers in individual i's municipality of residence m who are employed in the oil-sector in year $t-1$. As before, it is the interaction effect δ_{t+h} which is our object of interest. This coefficient tells us how much more non-oil households are affected by an oil price change occuring in year t if their municipality had a marginally higher $OilShare$ in year $t-1$. Note that this coefficient cannot be interpreted by itself, without also taking the estimate of γ into account. We therefore report the results from this exercise by computing the estimates' implied impulse responses to a 10 percent oil price change when the exposure variable, $OilShare$, is in the 5th and 95th percentile of its distribution. At the 5th percentile $OilShare = 0.03\%$, at the 95th percentile $OilShare = 7\%$.[27]

The estimated effects are displayed in Figure 8. Panels (a) to (c) plot the estimated responses of income, expenditure, and debt among non-oil households living in municipalities with a high (95th percentile) and low (5th percentile) share of oil-workers. We see that following an oil price hike of ten percent, income of non-oil households increases considerably more when the municipal oil exposure is higher. Non-oil income in the high-oil-share regions peaks two to three years after the oil price hike. Households in more exposed regions increase expenditure before their income rises, similar to what we previously found for the oil-sector households. The debt response, however, is similar in both type of municipalities, and close to zero. Overall, the findings indicate significant spillovers from the oil industry to the non-oil households, and that, as one would expect, the more oil intensive is the region, the larger are the spillovers.

[27]Note that we report estimated marginal responses to oil price fluctuations at different levels of oil shares, while taking into account the estimated coefficients on other covariates in our regression model. To arrive at the estimated responses reported in Figure 8, we set the value of these variables at their mean, and equal for all households. Alternatively, we could simply report the estimated interaction coefficient, and multiply it by the oil share.

Figure 8: Non-oil households' responses to a 10 % change in the oil price, in municipalities with high and low oil exposure.

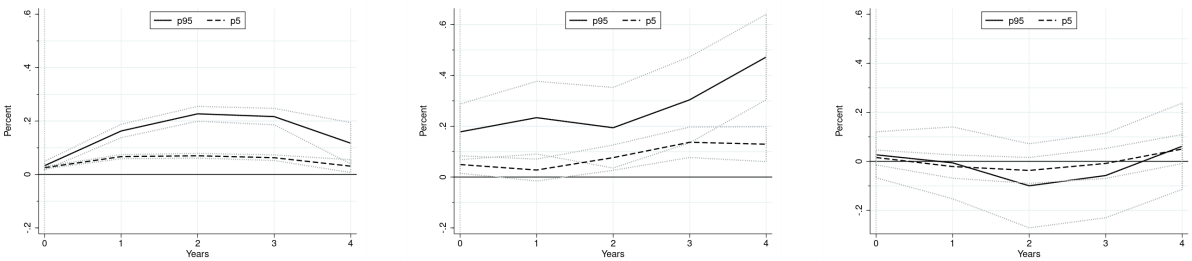

(a) Income (b) Expenditure (c) Debt

Notes: The responses of non-oil households' total after-tax income, expenditure and debt after a 10 percent increase in the oil price in period t, depending on the share of oil workers in their municipality of residence. The sample includes households classified as non-oil in $t-1$. Solid lines represent non-oil households residing in municipalities at the top 95th percentile of municipality-level oil-worker share ($\approx 7\%$), and dashed lines represent non-oil households residing in municipalities at the bottom 5th percentile of municipality-level oil-worker share ($\approx 0.3\%$). The plots follow from the estimates of Equation (3). Dotted gray lines are the 95 percent confidence bands. All variables are scaled by disposable income in $t-1$.

To facilitate interpretation of how big the estimated spillover effects are, we scale the income responses for non-oil households by the income response of households in the oil sector. In Table 4 we calculate the cumulated income responses at each horizon $h = 0$ to 4, relative to the total income response of households working in the oil sector. Because oil-sector households also experience income effects through spillovers, we use their estimated income response from a specificationwith time-fixed effects only, i.e. we intentionally do not control for local general equilibrium effects.

As already is clear from Figure 8, the spillover effects in regions with a small oil share (row 1) are smaller, compared with those in regions with a higher oil share (row 2). In the bottom 5th percentile, the cumulated 5-year income response of non-oil households amounts to 33 percent of that for oil households. In contrast, after 5 years non-oil households in oil-intensive regions have experienced a total income hike that is approximately equally high as it is for households working in the oil sector (96%).

One could in principle imagine a multitude of mechanisms behind these local general equilibrium effects. We find it useful to distinguish between two main channels. First, there might be a regional component in the labor demand from the oil sector itself, whereby after an oil price shock the oil sector to a greater extent increases its demand for labor in those localities where there are more oil-sector workers from before. One reason could be that workers are geographically immobile and prefer living close to their jobs, another reason could be that non-oil workers are closer substitutes to oil workers in oil-intensive regions. We term this the labor demand channel. Second, it could be that when the oil-

Table 4: Income effects for non-oil households by municipality-level exposure

	$\frac{\sum_{j=0}^{h} \Delta Y_{m,t+j}^{non-oil}}{\sum_{j=0}^{4} \Delta Y_{t+j}^{oil}}$				
	$h=0$	$h=1$	$h=2$	$h=3$	$h=4$
$5th\, percentile$	0.03	0.12	0.21	0.29	0.33
$95th\, percentile$	0.04	0.25	0.54	0.81	0.96

Notes: This table reports the implied income responses of non-oil households residing in municipalities with a high (p95) and low (p5) oil worker share, at each horizon, j, divided by the total income response of oil sector household over five years. We use the income response of oil sector households when we use time fixed effects in place of municipality-time fixed effects.

sector households increase their expenditure after an oil price hike, they buy locally produced goods and services, which in turn raises the income for local non-oil households. We term this the Keynesian demand channel. In sections 4.4 and 4.5, we assess how important each of these two channels are likely to be in accounting for our estimated local spillover effects.

4 Extensions

We extend our main analysis in certain key dimensions. First, we replace oil price changes with speculative oil price shocks as identified in Baumeister and Kilian (2016a). Second, we explore if intertemporal spending responses vary with initial deposit holdings and home equity. Third, we explore how results change if we narrow our analysis down to comparing oil-sector workers to non-oil workers in similar or different occupations. Fourth, we compute regional spillovers to households working in trade and non-trade industries separately.

4.1 Structurally identified oil price shocks

While oil price fluctuations are determined exogenously to each individual household's behavior, they are endogenous to global demand and supply factors. If these confounding factors affect oil sector households and non-oil sector households differently (i.e. if they correlate with our sector exposure variable), they might be driving our results above despite the controls for municipality-time fixed effects.

We therefore substitute oil price changes ($\Delta Oilprice\%_t$) with structurally identified oil price shocks in our econometric specification given in (2). The structurally identified oil price shocks are intended to net out systematic components from oil price fluctuations, in particular the endogenous responses of oil prices to global economic activity.

We use estimated shocks from the vector autoregression model presented in Baumeister and Kilian (2016a).[28] That model includes the real price of oil, global crude oil production, global real economic activity, and changes in the global stocks of crude oil. One source of oil price variation in this model is termed as "speculative demand shocks", which is distinguished from the endogenous responses of oil prices to the other three variables in the system as well as shocks to flow demand and flow supply. In order to combine this monthly shock series with our yearly administrative data, we aggregate the shock series to annual frequency as follows:[29]

$$\widehat{Oil}_t = \epsilon^{oil}_{tM12} + 2\epsilon^{oil}_{tM11} + ... + 12\epsilon^{oil}_{tM1} + 11\epsilon^{oil}_{(t-1)M12} + 10\epsilon^{oil}_{(t-1)M11} ... + \epsilon^{oil}_{(t-1)M2}, \tag{4}$$

where ϵ^{oil}_{tMj} denotes the VAR-identified shock in month j of year t. As seen, we compute a simple weighted sum of the shocks to arrive at the annual frequency. The weighting is motivated by the assumption that the oil price evolves as a unit root process, so that monthly shocks affect the current annual oil price change according to when they occured during the current and the previous year. Finally, the annualized shocks, $\widehat{Oil}_t$, are normalized such that a one unit increase represents a one standard deviation shock.

Figure 9 plots the original oil price series from the main section of our paper (solid line) together with the VAR-identified shocks (dashed line). The figure displays high co-movement between the two series, except from the end of the sample period, and that the VAR shocks are muted relative to the directly observed oil price changes. The high co-movement is not necessarily to be expected, since the VAR-shocks are purged of systematic correlation with past and present measures of world-wide economic activity.

[28] The VAR model presented in Baumeister and Kilian (2016a) in turn builds on the work by Kilian and Murphy (2014)

[29] The data is downloaded from https://www.aeaweb.org/articles?id=10.1257/jep.30.1.139 (Baumeister and Kilian (2016b))

Figure 9: Annual perecentage change in oil price and VAR-identified oil price shocks.

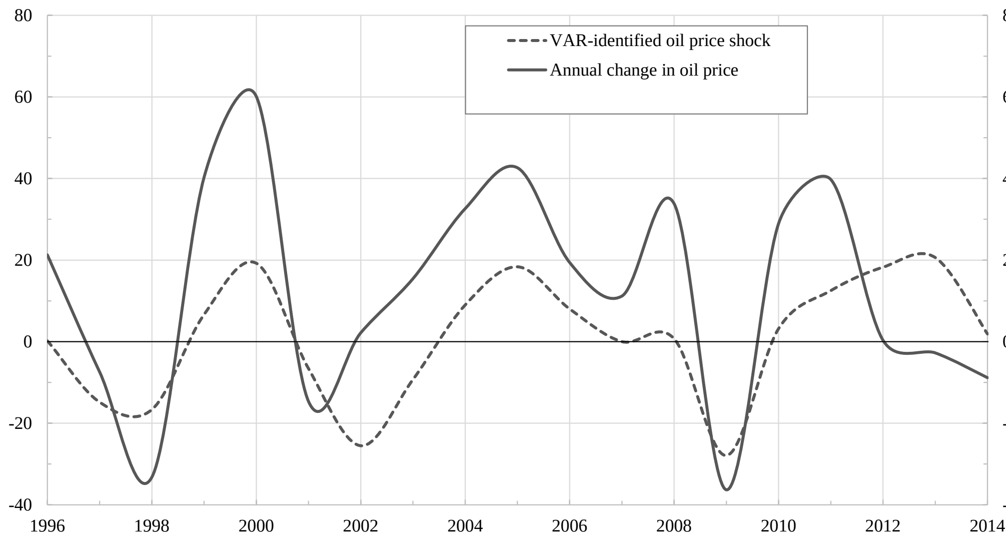

Notes: The figure plots the annual change in the price of oil (solid line), as used in specifications in all sections apart from this Section 4.1, and the annualized VAR-identified shocks from Baumeister and Kilian (2016a,b).

The regression model we now use corresponds to the baseline model in equation (2):

$$\begin{aligned}\frac{Y_{it+h} - Y_{it-1}}{Income_{it-1}} =& \alpha_i + \tau_{m,t} + \sum_{j=1}^{2} \boldsymbol{\beta}_j \mathbf{X}_{it-j} + \gamma_2^{EXP} EXP_{it-1} \\ &+ \delta_{t+h}^{EXP}(\widehat{Oil}_t \times EXP_{it-1}) + e_{it},\end{aligned} \tag{5}$$

where $\tau_{m,t}$ is the municipality-time fixed effect for municipality m at time t as before, and the oil price change ($\Delta Oilprice\%$) is replaced by the speculative demand shock ($\widehat{Oil}$).

Results are presented in Figure 10 as solid black lines, with 95 percent confidence bands around them. For comparison, we also plot our baseline point estimates from Section 3.1 in red. Overall, results are very similar in the two specifications, and our baseline point estimates mostly fall within the confidence bands of the new specification. Again, we observe an income response that is zero on impact and hump-shaped thereafter, coupled with an expenditure response that is positive on impact. As in our baseline estimates, expenditure falls below pre-shock levels in years three and four after the shock. Also as in the main specification, it is debt, and not existing deposits, that finances the anticipation effect in expenditure. The only substantive difference from the benchmark estimates, is the debt response which now drops below zero after the first period.

If anything our baseline responses are more muted than when we use the identified speculative demand shock series. Thus, it does not seem that endogeneity of oil price fluctuations to global economic activity or oil inventories is an important concern for interpreting our baseline estimates that use the directly observed oil price change.

Figure 10: Responses to speculation-induced oil price shocks identified in a structural VAR.

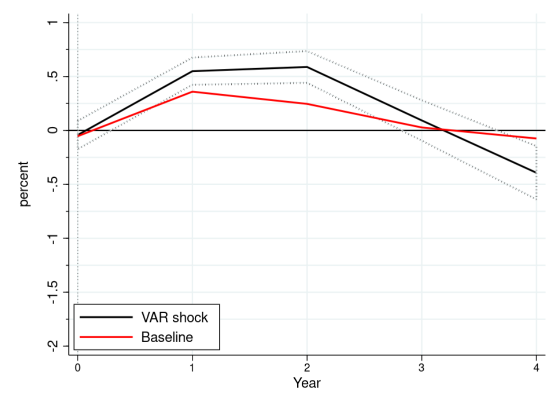

(a) Income

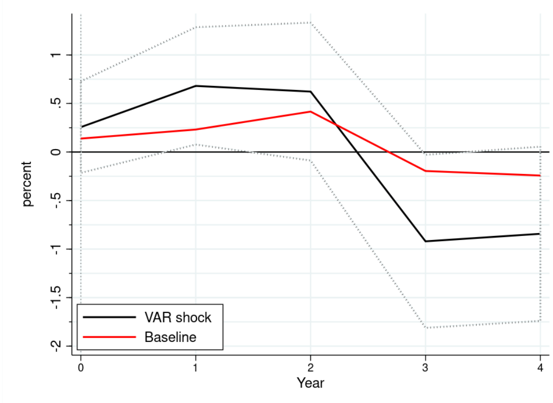

(b) Expenditure

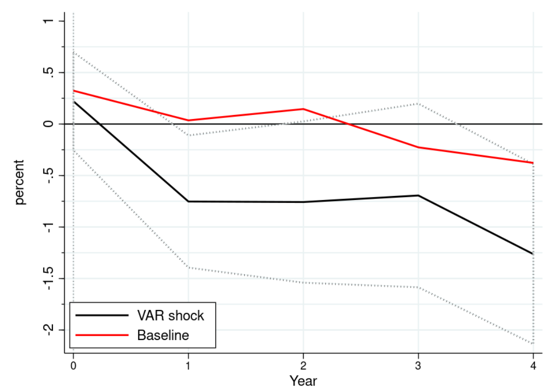

(c) Debt

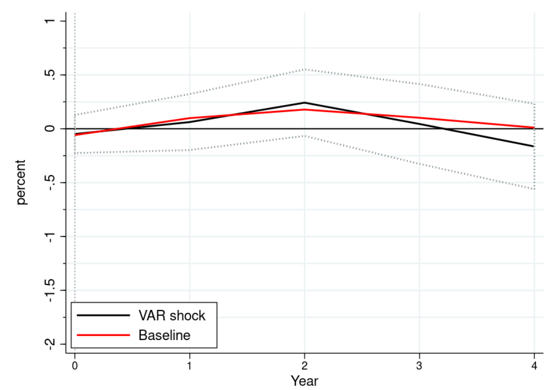

(d) Deposits

Notes: The black solid lines plots the estimated coefficients (δ^k_{t+h} in equation (5)) using the VAR-identified speculative demand shocks from Baumeister and Kilian (2016a) aggregated to annual frequency (see text for details). Dotted lines represent the 95 percent confidence bands of these point estimates. Red lines plot the baseline results from specification (2) as previously displayed in Figure 4 and Figure 5. All responses are those of oil households relative to non-oil households after a one standard deviation shock in time $t = 0$. Disposable income is defined as the sum of labor income, transfers and capital income, after tax. Expenditure is imputed from the household budget constraint (see e.g. Fagereng and Halvorsen (2017)). All variables are scaled by disposable income in $t-1$. The classification of sectors is based on Statistics Norway (2017), see Section 2 for details.

4.2 Do anticipation effects vary with initial deposits or home equity?

In Section 3.1 we found that on average across households, the initial expenditure response is financed by increased borrowing rather than by decumulation of deposits. We now go one step further in this direction and explore how the intertemporal expenditure responses vary with households' initial deposit and home equity holdings. We estimate the following relation separately for different groups g, where each group is defined by initial (time $t-1$) holdings of deposits or home equity

$$\begin{aligned}\frac{Y_{it+h}-Y_{it-1}}{Income_{it-1}} =& \alpha_{i,g}+\tau_{t,m,g}+\sum_{j=1}^{2}\beta_{j,g}\mathbf{X}_{it-j}+\gamma_{2,g}EXP_{it-1} \\ &+\delta_{gt+h}^{EXP}(\Delta Oilprice\%_t \times EXP_{it-1})+e_{it},\ \forall i\in g. \end{aligned} \tag{6}$$

Splitting our sample into four groups reduces the statistical power in our regressions. We therefore pool together the oil households and oil service households into one group, such that EXP_{it-1} is a dummy variable separating these two groups from non-oil households.

This exercise is motivated by the recent emphasis on liquidity constraints in the macroeconomic literature (see for instance Kaplan and Violante (2014)). Within this recent modelling tradition, a significant share of wealthy households can be liquidity constrained because they primarily hold illiquid assets which cannot readily be converted to consumption. Such households then behave in a hand-mouth manner where consumption tracks income. [30]

By this logic, we would expect that households with low deposits are more constrained and hence less inclined to raise their expenditure immediately in response to anticipated future income hikes. However, when we split households by deposit holdings, we do not find such a pattern. Figure 11b and Figure 11c plots the income and expenditure trajectories among households that differ by their initial deposit holdings, namely the bottom (red lines) and top (blue lines) deposit quartiles. The income response of the two groups are quite similar, and in both groups we see a similar anticipation effect in expenditure as before.[31] Expenditure of the low-deposits households seem to respond more relative to their income response than the high-deposits households do. This is also apparent in Figure 11a, which is based on the same set of regression coefficients. Here we calculate and graph the initial response of expenditure, as well as debt and deposits, relative to the total income response over the entire five-year horizon. We observe that among households with low deposits,

[30] This emphasis on asset liquidity is in turn largely motivated by the fact that in the data, households tend to hold a relatively small share of their wealth in deposits, and a need to size up household's marginal propensity to consume so as to align with microeconomic estimates.

[31] Impulse responses with confidence bands for all four quartiles and for the variables income, expenditure, debt and deposits is reported in the Appendix, Figure 22.

expenditure increases by roughly the same amount as the total income response on impact, and this is almost entirely financed by debt. High-deposit households increase expenditure less relative to their own total income rise, and partly finance this by running down on their deposits.

Figure 11: Responses by quartiles of deposits in t-1.

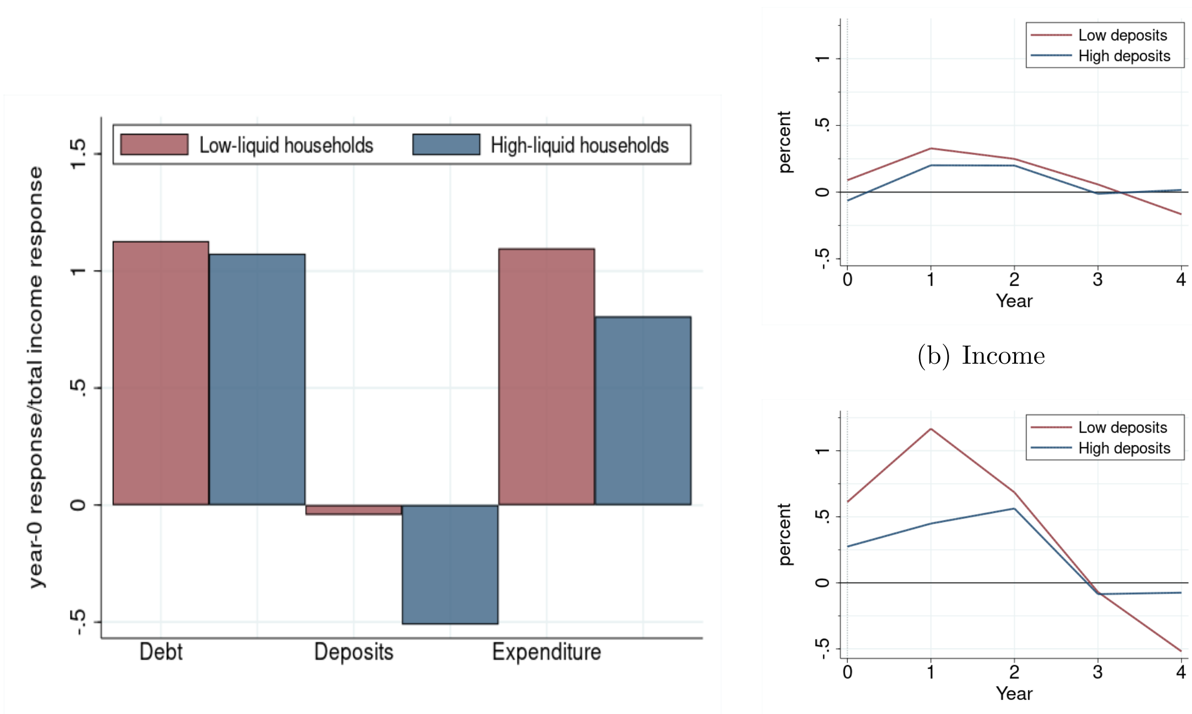

(a) Response in year-0, relative to total 5-year income response

(b) Income

(c) Expenditure

Notes: The figures follow from the estimated coefficients δ^{EXP}_{t+h} *in Equation (6). Oil and oil service sectors are pooled, such that the sector variable is a dummy variable equal to one if the household earns income from the oil or oil-service sector, and zero otherwise. Households are divided into quartiles of deposit holdings in year* $t-1$*. Figure (a) plots the responses of household balance sheet components for households in the lowest quartile of holdings of deposits (red bars) and the highest quartile (blue bars) in period zero, relative to the total income response over the 5 years after a ten percent increase in the oil price. Plots (b) and (c) show the impulse responses of income and expenditure for each of the two groups. Red lines represent households with low deposits and blue lines represent households with high deposits. Disposable income is defined as the sum of labor income, transfers and capital income, after tax. Expenditure is imputed from the household budget constraint (see e.g. Fagereng and Halvorsen (2017)). All variables are scaled by disposable income in* $t-1$*. The classification of sectors is based on Statistics Norway (2017), see Section 2 for details.*

We also split households by their initial home equity to investigate whether mortgage market credit constraints curb the anticipation effect in expenditures among highly indebted households. Our results from a variety of exercises in this spirit are inconclusive. Hence, while our baseline estimates indicate that borrowing is important for households to respond in anticipation of income hikes, our heterogeneity estimates do *not* suggest that low home equity prevents such expenditure responses. We report the results from such exercises in Figure 23 in the appendix.

4.3 Control groups: oil occupations versus non-oil occupations

Our empirical strategy relies on our groups being differently exposed to oil price fluctuations via their jobs, but it also relies on the assumption that non-oil households constitute a relevant control group. As an extension to our main analysis, we therefore explore the role of the control group further by using data on occupational codes of individuals to provide a more fine-grained analysis of the non-oil control group. Our goal is to construct more or less similar control group. For instance, we want to compare responses of engineers working in the oil sector to engineers working outside the oil sector in one exercise, and to, say, teachers, in another exercise. If the labor market is competitive, we would expect that the more similar is the occupation of the non-oil household control group to the oil-sector workers, the smaller will the differential income, and therefore expenditure, responses be. If the labor market is less competitive, then income responses may differ across households within the same occupation, in which case we can investigate differential responses of other variables of interest too.

We use data on occupations that are available after 2005 only. We conduct a novel classification of all occupation codes into "oil occupations" and "non-oil occupations", irrespective of sector (NACE code), but depending on the share of oil sector individuals working within each occupation category. Our procedure is as follows: we first calculate the share of individuals working in the oil sector or the oil service sector in relatively broad occupational categories. There are roughly 100 such categories.[32] We then classify an occupation code as an oil occupation if the share of oil-sector individuals with this occupation code exceeds one percent of the total number of individuals in that occupation. More precisely, we classify each occupation group with the following categorical variable:

$$OilOccupationDummy = \begin{cases} 1, & \text{if } \frac{\sum OilSectorWorkers}{\sum Workers} > 0.01 \\ 0, & \text{otherwise} \end{cases}$$

[32] The data is described in Statistics Norway (1998), available at https://www.ssb.no/arbeid-og-lonn/yrkeskatalogen. To give an idea of the level of aggregation, some examples of these categories are: "Physical, mathematical and engineering science professionals", another is "Engineering science associate professionals", and a third is "Extraction and building trades workers." Ten out of the 100 categories are occupations categorized as "'Legislators, senior officials and managers." In the subcategories of managers, many individuals are classified as working in the oil sector, according to our classification. However, since these categories do not give any description of type of education, or type of work, we choose not to include these individuals in our classification of occupations.

The highest share of oil individuals in any occupation is 21 percent ("Prosessoperatorer mv."). A total of 9 out of the 90 occupations considered are coded as oil occupations.[33] 85 percent of the individuals classified as working in the oil sector fall within these nine oil occupation categories. At the same time, only six percent of the individuals that work in an oil occupation, work in the oil sector.

We rerun specification 2, but replace the full non-oil household sample as the control group. First we replace it with the oil occupations group ($OilOccupationDummy = 1$), and thereafter we use the non-oil occupations group ($OilOccupationDummy = 0$). The income and expenditure responses from these regressions are displayed in Figure 12.[34] Blue lines represent point estimates when we only use non-oil households with oil occupations as the reference group, and red lines represent point estimates when we only use non-oil households with non-oil occupations as the reference group. We also plot our baseline point estimates from Section 3.1 for comparison.

Figure 12: Income and expenditure responses of oil households relative to three different control groups.

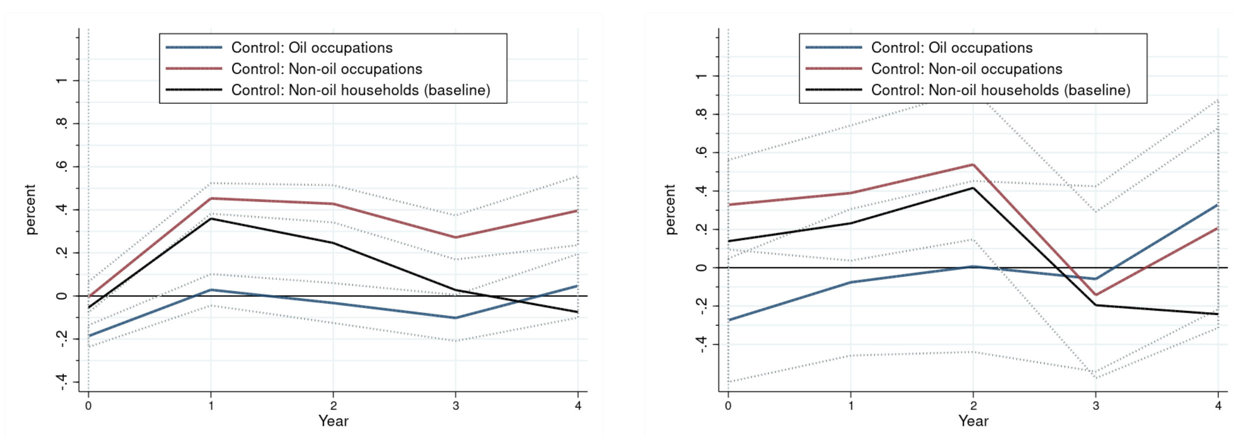

(a) Income (b) Expenditure

Notes: The figure plots the responses of household income and expenditure for oil households relative to non-oil households in every period $t+h$ relative to income in period $t-1$, after a 10 percent increase in the oil price in period t. The plots follow from the estimated coefficients δ_{t+h}^{EXP} as in Equation (2), with $EXP_{t-1} = sector_{t-1}$ as exposure variable, but where the reference group of non-oil households is altered to include only non-oil households with oil occupations (blue lines) or non-oil households with non-oil occupations (red lines). Dotted lines represent the 95 percent confidence bands of these point estimates. Grey lines plot the baseline responses where all non-oil households are included in the reference group, as previously reported in Figure 4. See text in this section for details on the classification of oil and non-oil occupations. Disposable income is defined as the sum of labor income, transfers and capital income, after tax. Expenditure is imputed from the household budget constraint (see e.g. Fagereng and Halvorsen (2017)). All variables are scaled by disposable income in $t-1$. The classification of sectors is based on Statistics Norway (2017), see Section 2 for details.

[33]The nine categories are listed in the appendix, see section A.6.

[34]Results for the oil service sector are reported in the appendix, see Figure 24

Panel (a) shows that within occupations (blue line), the differential income response is close to (and not significantly different from) zero, except in the impact year where oil-sector households' income drops somewhat relative to that of non-oil households. The corresponding within-occupation differential expenditure response displayed in panel (b) is similar to that of the income response, and it is not statistically different from zero in any period. In contrast, if our control group contains only households in non-oil occupations, we see from panel (a) that the differential of income response widens, and so does the expenditure response in panel (b).

These results show that for our purposes it is problematic to infer anticipation effects in expenditure or spousal labor market insurance effects within occupation categories. The reason is, if the labor market is competitive, an oil price induced hike of labor demand in the oil sector will lift the earnings of both oil and non-oil workers within oil occupations. This could happen both because non-oil workers switch to the oil sector (extensive margin), or because non-oil firms are forced to increase salaries in order to retain workers (intensive margin). On the other hand, when we compare with occupations less likely to be hired in the oil sector, the differential earnings response becomes greater and it becomes possible to trace out differential responses along the other margins too. Because most households work in non-oil occupations, these dominate the control group in our baseline analysis and thus our baseline estimates are similar to what we find when using only non-oil occupations in our control group.

4.4 Regional exposure: traded versus non-traded industries

The results in Section 3.3 suggest that oil price fluctuations spill over to the income of non-oil households. There are two main channels for such spillovers that are natural to consider. First, they might be due to a regional component in the labor demand from the oil-sector itself, whereby the oil sector to a greater extent increases its demand for labor in those localities where there are more oil-sector workers from before. Second, the spillovers might be due to the Keynesian demand channel, whereby greater spending by oil workers raises the income for other households too. In this and the next section, we attempt to shed light on how important each of these two channels is for explaining our regional spillover estimates.

If the Keynesian demand channel is important, then non-traded sectors should be affected more than traded sectors, as the latter rely on global rather than local demand. To investigate this hypothesis, we follow Mian and Sufi (2014), and split the non-oil sector into traded, non-traded and "other" industries. Then we use the same type of regressions as before, but with an added interaction term for the industry, creating a three-way interaction term of oil price change, oil share in the municipality and industry. That is, in model (3) we add the variable $\Delta Oilprice\%_t \times OilShare_{imt-1} \times Industry_{it-1}$, as well as the two-way interaction terms $Industry_{t-1} \times OilShare_{imt-1}$ and $Industry_{t-1} \times \Delta Oilprice\%_t$ and the $Industry_{t-1}$ variable itself.

The results are shown in Figure 13, where we plot the difference between households living in high-oil-share municipalities and low-oil-share municipalities on the y-axis, as in our main analysis in Section 3.3. Three lines represent the effects for households in the traded industry (dashed), in the non-traded industry (dotted), and in industries that are neither classified as traded nor non-traded (other industry, solid line). To be classified as a traded-industry household, at least one member of the household earns income from the traded industry while none earn income from the non-traded industry or the oil sector. Vice versa, non-traded households have at least one income from the non-traded industry and no income from the traded industry or the oil sector. The remaining households without income from the oil sector are classified as belonging to "other industries."

Figure 13: The effect of regional exposure on non-oil households' responses to a 10 % change in the oil price, by industry of employment.

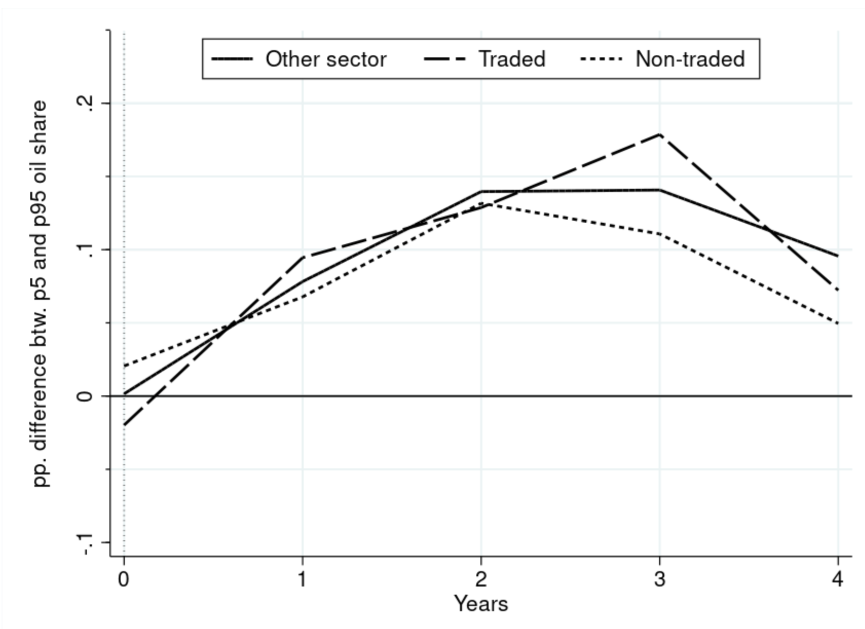

*Notes: Each line displays how much more non-oil households are affected by an oil price change of 10 percent if they reside in a municipality at the top 95th percentile rather than at the 5th percentile of the local exposure distribution. The dashed and dotted lines are for non-oil households working in industries classified as producing tradables and non-tradables, respectively, while the solid curve is for other non-oil households. See text for details on categorization. The plots follow from the estimated coefficients δ^{EXP}_{t+h} in Equation (2) but with time-fixed effects instead of municipality*time-fixed effects, and an added interaction term for industry.*

Our estimated income responses across households in the three industries, which indicates that the aggregate demand channel is *not* the main driving force behind our baseline regional estimates. Alternatively, if aggregate demand is a driving force behind responses in the non-traded industries, local labor demand effects affect workers in the traded industries to the same extent. It is possible that aggregate demand for goods and services primarily benefit workers in the non-traded industries, whereas oil-sector demand for labor to a greater extent benefits workers in traded industries. The latter is supported by the occupations data we used in Section 4.3. 59% out of all traded-industry households are in oil occupations, while only 47% of non-traded industry households are in oil occupations. And, importantly, the share of oil-occupation workers in municipalities' traded industry is increasing with municipal oil share. If we stratify municipalities by their share of oil-sector households, we see that in the lower quartile of municipal oil share, oil occupations have a 48% fraction in the traded industry, whereas in the upper quartile the same fraction is 61%.

As discussed in Sections 2.3 and 3.3, a concern with our baseline spill-over estimates is that misclassification error (that some oil households are classified as non-oil households) is positively correlated with municipalities' oil intensity. This will induce an upward bias in our spill-over estimates because they will also include the effect on these misclassified non-oil, but in reality directly oil-exposed, households.[35]

To address this concern, we therefore study households where such misclassification error is likely to be small and independent of municipal oil share. Households earning income from the retail industry, we argue, is such a group.[36] Hence, we split the broader non-traded household group into its more fine-grained sub-groups. We then re-run regression (3), this time with separate regressions for the non-oil households classified as working in the service, retail, and construction industries.

Our results are reported in Figure 14. Here, we plot the implied impulse responses for non-oil households residing in municipalities at the top 95th percentile of municipality-level oil-worker share for different sub-groups within the non-traded industry (black lines). We also plot the baseline coefficients (as previously reported in Figure 8a) where all non-oil households are included (red line). The figure shows some differences between the groups, mostly with respect to timing. Point estimates from our baseline tends to fall within the confidence bands of our new estimates. Figure (b) shows that there are spill-over effects even in the retail industry group, where misclassification is likely small, and independent of oil intensity in the region. Hence, our qualitative results on spill-overs in Section 3.3 do not appear to be driven by misclassification of households across sectors.

[35]Recall that the opposite is true for our sector exposure estimates. The same misclassification error will create a *negative* bias in these estimates.

[36]Workers in the traded industries on the other hand, are more likely to have a higher correlation with the misclassification error, which in turn might explain why we don't observe any difference between the traded and non-traded non-oil households above, Figure 13.

Figure 14: Non-oil households' responses to a 10 % change in the oil price. By industry of employment.

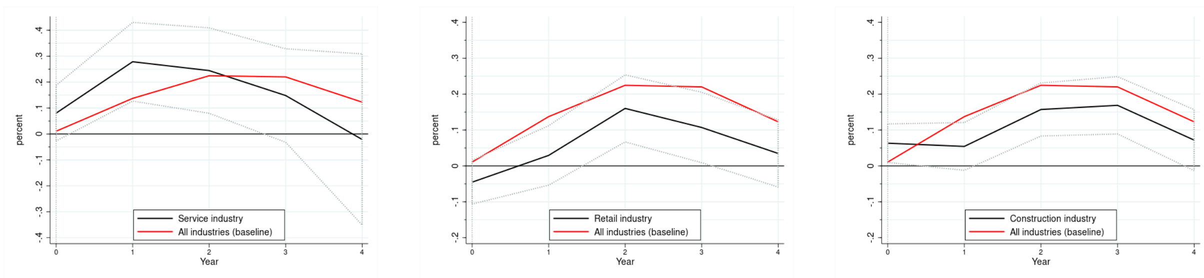

(a) Service industry (b) Retail industriy (c) Construction industriy

Notes: The figure plots the income responses after a 10 percent increase in the oil price in period t, for non-oil households residing in municipalities at the top 95th percentile of municipality-level oil-worker share, split by the households' industry of employment. The plots follow from the estimates of Equation (3), run separately for each industry sub category. Black lines represent the response of non-oil households earning income from the service industry (panel a), the retail industry (panel b) and the construction industry (panel c), respectively. Dotted lines represent the 95 percent confidence bands of these point estimates. Red lines represent the baseline response, as previously reported in Figure 8a, i.e. the response of all non-oil households (irrespective of industry) residing in municipalities at the top 95th percentile of municipality-level oil-worker share. The oil-worker share at the top 95th percentile is approximately 7 percent.

To sum up, the similarity of income responses for households in traded and non-traded industries imply that Keynesian-style aggregate demand effects at most play a limited role in explaining the local general equilibrium effects estimated in Section 3.3. Instead it is likely that the direct labor demand channel, through which oil-sector demand for labor stimulates incomes in oil occupations, is the main driver behind the spillovers we estimated.

4.5 Regional exposure: occupations within industries

We here go one step further to assess the importance of the direct labor demand channel versus the Keynesian aggregate demand channel, in explaining the regional spillover effects to non-oil workers. To this end, we build on the preceding sections and split households by occupation and industry. The underlying motivation is that if the spillover effects in Section 3.3 are due to labor demand, we should see greater income responses for workers in oil occupations. In contrast, a Keynesian aggregate demand effect should impact income similarly for oil and non-oil occupations.

Again the econometric specification is a version of Equation (3). The coefficient of interest is δ_{t+h}, which says how much stronger a non-oil household's income responds to an oil price shock if their municipality is inhabited by marginally more oil-sector workers. We estimate this coefficient for non-oil and oil occupations separately. First, we do this across all non-oil households. Then we split the non-oil sector into traded and non-traded industries as in

Section 4.4, and estimate Equation (3) for oil and non-oil occupations.

The results are exhibited in Figure 15. Panel (a) shows that the regional spillover estimates across non-oil industries are larger for households in oil occupations. This is consistent with the direct labor demand channel being important. In contrast, if the Keynesian aggregate demand channel lay behind our estimated spillovers, we would instead expect the two response curves to behave similarly. In panel (b), we see that among households in the traded industries, income responds most for those who are in oil occupations. Again, this is consistent with the labor demand channel being important. In panel (c) we see less of such a difference in the non-traded industry, but still the same qualitative pattern. Moreover, if we compare the non-oil income responses in traded and non-traded industries, they are quite similar. In contrast, if the Keynesian demand channel was the main driver of spillovers, and there were some short-run segmentation between the two sectors, then we would expect income to respond more for non-oil occupations in the non-traded than in the traded industries. That does not appear to be the case.

Figure 15: The effects of municipal oil exposure on non-oil households' responses across occupations and industries.

(a) All industries (b) Traded industries (c) Non-traded industries

Notes: The figure plots the income responses after a 10 percent increase in the oil price in period t, for non-oil households residing in municipalities at the top 95th percentile of municipality-level oil-worker share, split by the households' industry of employment and occupation. The plots follow from the estimates of Equation (3), run separately for each sub group. In all panels red lines represent the response of households in oil occupations, and blue lines represent the response of households in non-oil occupations. Panel a pools all industries together, whereas panel b and panel c split households further into groups for the traded and non-traded industries, respectively. Dotted lines represent the 95 percent confidence bands of the point estimates. The oil-worker share at the top 95th percentile is approximately 7 percent. Definitions and categorization of occupation and industry types are provided in the text, see Section 4.3 and Section 4.4.

Taken together, this and the previous section give a quite consistent picture when it comes to the relative importance of the direct labor demand channel versus the Keynesian aggregate demand channel. It is the former that seems to be the most important to explain why non-oil households benefit more from oil price hikes in those municipalities with more oile-sector workers.

5 Theoretical interpretation

Our expenditure estimates in Section 3.1 indicate anticipation effects and a tendency for "overshooting", in the sense that expenditure first responds positively and more than income, and thereafter drops below zero in years three and four. A likely explanation is durable expenditures. In the presence of durables, a household who receives news of a short-run income hike will want to purchase more durable goods immediately. Moreover, if durables goods are indivisible or if buying them involves fixed transaction costs, then households will shift planned future purchases forward in time when they observe higher income. Our imputed expenditure measure includes durables, and so these mechanisms seem plausible explanations of our overshooting estimates in Section 3.1.

Notably, purchases of durable goods might also explain why Keynesian multiplier effects seem of limited importance for the regional general equilibrium effects estimated in Section 3.3. The reason is, durable goods typically involve less local inputs than non-durables, in particular services.

Finally, the spousal earnings responses are relatively high in the medium run, before they seem to disappear. In a simple model of spousal labor supply, we would instead see a smoother wealth effect, where the spouse cuts his or her response permanently. The joint pattern of a high short-term expenditure response and a relatively big spousal earning contraction, could indicate that in the household welfare function there is complementarity between the spouses' leisure and the households' durable goods consumption.

6 Conclusion

We offer micro estimates of how oil price movements affect households in an oil exporting economy along key margins. First, we show that an oil price hike leads to a gradual, hump-shaped increase in the income of households employed in the oil sector. In contrast, household expenditure reponds immediately, consistent with an anticipation effect in consumption. Second, we estimate that within households with one worker in the oil sector, spousal earnings responses are sizable and thus an important margin for household-level income insurance in the face of oil price fluctuations. Third, we estimate substantial general equilibrium spillover effects to households who are not employed in the oil sector before the oil price increases, as non-oil household incomes increase significantly more following an oil price hike in municipalities where there are more oil-sector workers. Moreover, corroborative evidence across industries and occupations indicates that these general equilibrium effects primarily stem from the oil-sector's increased labor demand rather than Keynesian demand multipliers effects.

To advance our understanding of macroeconomic shock propagation, an ambition of current research is to develop models that are consistent with micro as well as macro evidence. Our study contributes to this agenda, with estimates that are particularly relevant for how export prices transmit to small open economies. At the partial equilibrium level, our evidence point toward models where anticipated income shocks shift expenditure forward in time, as in theories of durable consumption, and where spousal labor supply provides partial insurance against the idiosyncratic risk caused by export price fluctuations. At the general equilibrium level, our results point toward models where competitive labor markets cause spillovers from export price movements to households employed in sectors that are not directly exposed to the price changes, but not toward models where aggregate demand multipliers are particularly important.

References

Acemoglu, D., A. Finkelstein, and M. J. Notowidigdo (2013): "Income and health spending: Evidence from oil price shocks," *Review of Economics and Statistics*, 95, 1079–1095.

Agarwal, S. and W. Qian (2014): "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore," *American Economic Review*, 104, 4205–30.

Akram, Q. F. (2004): "Oil prices and exchange rates: Norwegian evidence," *The Econometrics Journal*, 7, 476–504.

Attanasio, O., H. Low, and V. Sánchez-Marcos (2005): "Female labor supply as insurance against idiosyncratic risk," *Journal of the European Economic Association*, 3, 755–764.

Auclert, A., B. Bardoczy, and M. Rognlie (2020a): "MPCs, MPEs and Multipliers: A Trilemma for New Keynesian Models," NBER Working Papers 27486, National Bureau of Economic Research, Inc.

Auclert, A., M. Rognlie, M. Souchier, and L. Straub (2021): "Exchange Rates and Monetary Policy with Heterogeneous Agents: Sizing up the Real Income Channel," Working paper.

Auclert, A., M. Rognlie, and L. Straub (2018): "The Intertemporal Keynesian Cross," Tech. rep., National Bureau of Economic Research.

——— (2020b): "Micro Jumps, Macro humps: Monetary Policy and Business Cycles in an Estimated HANK Model," Tech. rep., National Bureau of Economic Research.

Baker, S. R., L. Kueng, S. Meyer, and M. Pagel (2021): "Consumption Imputation Errors in Administrative Data," *The Review of Financial Studies*, hhab087.

Bardóczy, B. (2020): "Spousal Insurance and the Amplification of Business Cycles," *Unpublished Manuscript, Northwestern University.*

Baumeister, C. and J. D. Hamilton (2019): "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks," *American Economic Review*, 109, 1873–1910.

Baumeister, C. and L. Kilian (2016a): "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us," *Journal of Economic Perspectives*, 30, 139–60.

——— (2016b): "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us," *Journal of Economic Perspectives*, 30, Data Set.

Bayer, C., B. Born, and R. Luetticke (2020): "Shocks, Frictions, and Inequality in US Business Cycles," CEPR Discussion Papers 14364.

Bergholt, D., V. H. Larsen, and M. Seneca (2019): "Business cycles in an oil economy," *Journal of International Money and Finance*, 96, 283–303.

Bjørnland, H. (2018): "Fiscal Management of Oil Resources in Booms and Busts," *Policy Perspective*, 1–7.

Bjørnland, H. C. and L. A. Thorsrud (2016): "Boom or gloom? Examining the Dutch disease in two-speed economies," *The Economic Journal*, 126, 2219–2256.

Bjørnland, H. C., L. A. Thorsrud, and R. Torvik (2019): "Dutch disease dynamics reconsidered," *European economic review*, 119, 411–433.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016): "Consumption Inequality and Family Labor Supply," *American Economic Review*, 106, 387–435.

Browning, M. and S. Leth-Petersen (2003): "Imputing Consumption from Income and Wealth Information," *The Economic Journal*, 113, F282–F301.

Charles, K. K., Y. Li, and M. Stephens Jr (2018): "Disability benefit take-up and local labor market conditions," *Review of Economics and Statistics*, 100, 416–423.

Chodorow-Reich, G., P. T. Nenov, and A. Simsek (2021): "Stock Market Wealth and the Real Economy: A Local Labor Market Approach," *American Economic Review*, 111, 1613–57.

Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017): "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging," *American Economic Review*, 107, 3550–88.

Dornbusch, R. (1976): "The Theory of Flexible Exchange Rate Regimes and Macroeconomic Policy," *The Scandinavian Journal of Economics*, 78, 255–275.

Dornbusch, R. and S. Fischer (2017): *Macroeconomics*, McGraw Hill, 13 ed.

Druedahl, J., E. B. Jensen, and S. Leth-Pedersen (2018): "The Intertemporal Marginal Propensity to Consume out of Future Persistent Cash-Flows," Tech. rep.

Eika, L., M. Mogstad, and O. L. Vestad (2020): "What can we learn about household consumption expenditure from data on income and assets?" *Journal of Public Economics*, 189, 104163.

Fagereng, A. and E. Halvorsen (2017): "Imputing consumption from Norwegian income and wealth registry data," *Journal of Economic and Social Measurement*, 42, 67–100.

Fagereng, A., M. B. Holm, and G. J. Natvik (2021): "MPC Heterogeneity and Household Balance Sheets," *American Economic Journal: Macroeconomics*, 13, 1–54.

Fleming, J. M. (1962): "Domestic financial policies under fixed and under floating exchange rates," *Staff Papers*, 9, 369–380.

Galí, J. and T. Monacelli (2005): "Monetary Policy and Exchange Rate Volatility in a Small Open Economy," *Review of Economic Studies*, 72, 707–734.

Halla, M., J. Schmieder, and A. Weber (2020): "Job displacement, family dynamics, and spousal labor supply," *American Economic Journal: Applied Economics*, 12, 253–87.

Hamilton, J. D. (1983): "Oil and the macroeconomy since World War II," *Journal of political economy*, 91, 228–248.

Holm, M. B., P. Paul, and A. Tischbirek (2021): "The transmission of monetary policy under the microscope," *Journal of Political Economy*, 129, 2861–2904.

Jordà, Ò. (2005): "Estimation and inference of impulse responses by local projections," *American economic review*, 95, 161–182.

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

Kaplan, G. and G. L. Violante (2014): "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 82, 1199–1239.

Keane, M. P. and E. S. Prasad (1996): "The employment and wage effects of oil price changes: a sectoral analysis," *The Review of Economics and Statistics*, 389–400.

Kilian, L. (2009): "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99, 1053–69.

Kilian, L. and D. P. Murphy (2014): "The role of inventories and speculative trading in the global market for crude oil," *Journal of Applied econometrics*, 29, 454–478.

Kline, P. (2008): "Understanding sectoral labor market dynamics: An equilibrium analysis of the oil and gas field services industry," .

Kolsrud, J., C. Landais, and J. Spinnewijn (2020): "The value of registry data for consumption analysis: An application to health shocks," *Journal of Public Economics*, 189.

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

Mankart, J. and R. Oikonomou (2016): "The rise of the added worker effect," *Economics Letters*, 143, 48–51.

Mian, A. and A. Sufi (2014): "What explains the 2007–2009 drop in employment?" *Econometrica*, 82, 2197–2223.

Mundell, R. A. (1963): "Capital mobility and stabilization policy under fixed and flexible exchange rates," *Canadian Journal of Economics and Political Science/Revue canadienne de economiques et science politique*, 29, 475–485.

Obstfeld, M. and K. Rogoff (1995): "Exchange Rate Dynamics Redux," *Journal of Political Economy*, 103, 624–660.

Olsson, J. et al. (2019): "Structural transformation of the labor market and the aggregate economy," *Unpublished Manuscript, University of Amsterdam.*

Pruitt, S. and N. Turner (2020): "Earnings Risk in the Household: Evidence from Millions of US Tax Returns," *American Economic Review: Insights*, 2, 237–54.

Statistics Norway (1998): "Standard classification of occupations," Statistics norway Report.

——— (2017): "Employment in the petroleum industry and related industries 2016," Statistics norway Report.

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

A Appendix

A.1 Pre-trends

Figures 16 and 17 plots the same results as in the main text, but also the response of households in years *before* the oil price hike. That is, in the Figures below, in the years before 0, we plot the estimated response of oil households to a future oil price hike in year 0.

Figure 16: Responses to a 10 % change in the oil price. Oil households relative to non-oil households.

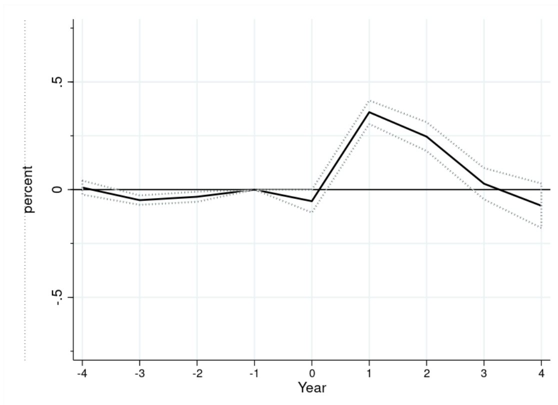

(a) Income

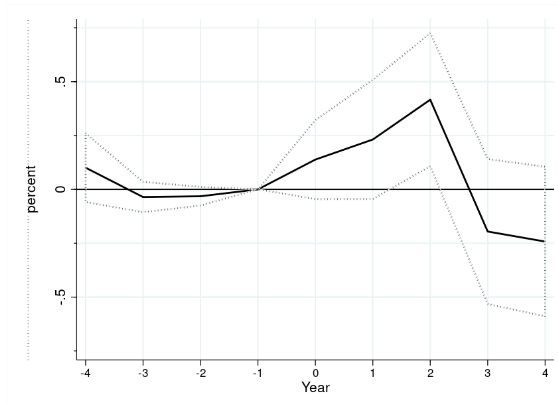

(b) Expenditure

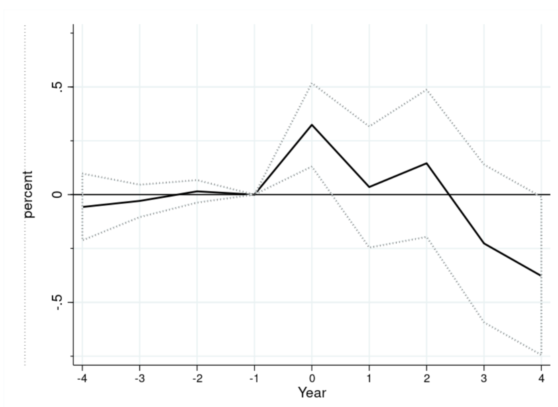

(c) Debt

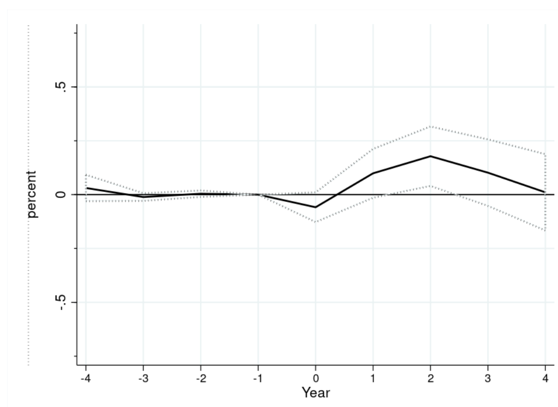

(d) Deposits

Notes: The figure plots the balance sheet responses of oil households, relative to non-oil household, before and after after a 10 percent increase in the oil price in period t. The plots follow from the estimated coefficients δ^{EXP}_{t+h} in Equation (2), with $EXP_{t-1} = sector_{t-1}$ as exposure variable. Disposable income is defined as the sum of labor income, transfers and capital income, after tax. Expenditure is imputed from the household budget constraint as explained in the main text. All variables are scaled by disposable income in $t-1$. The classification of sectors is based on Statistics Norway (2017), see Section 2 for details.

Figure 17: Responses to a 10 % change in the oil price. Oil-service households relative to non-oil households.

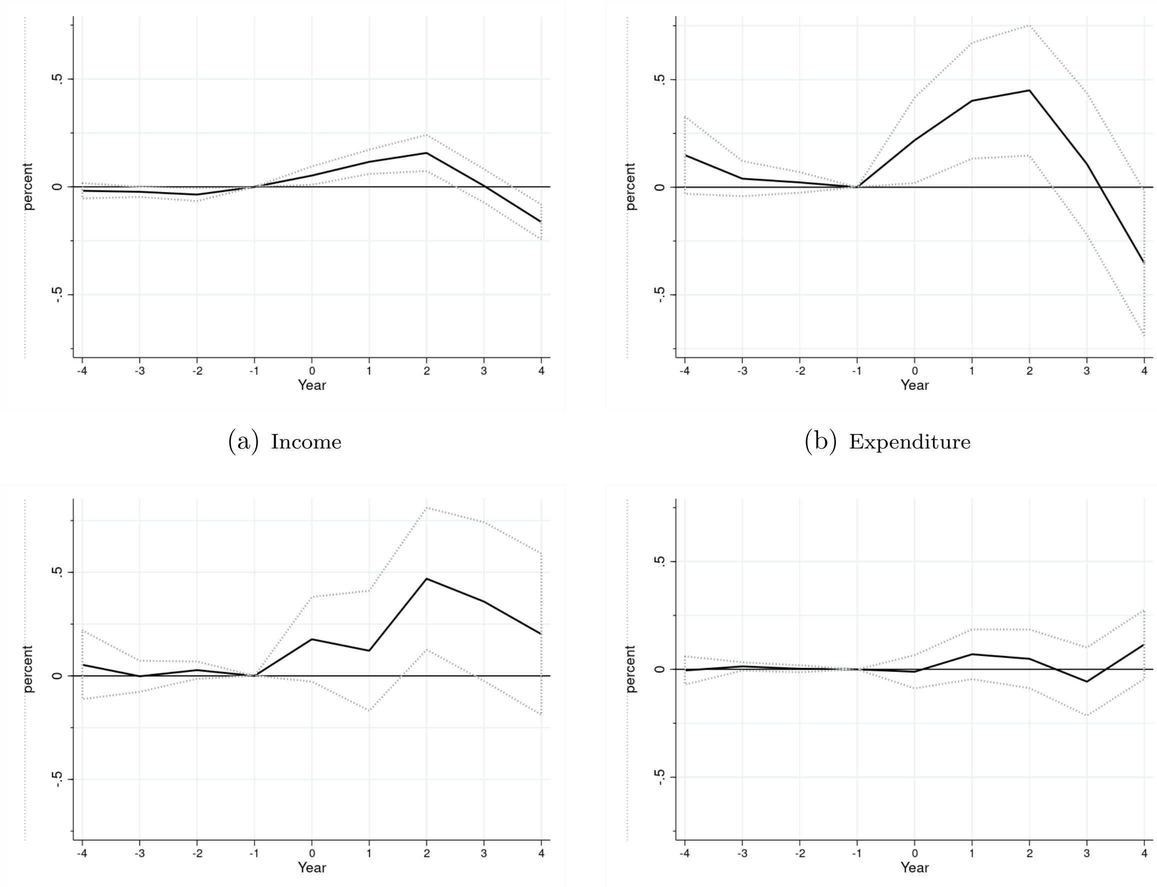

(a) Income

(b) Expenditure

(c) Debt

(d) Deposits

Notes: The figure plots the balance sheet responses of oil service households, relative to non-oil household, before and after after a 10 percent increase in the oil price in period t. The plots follow from the estimated coefficients δ_{t+h}^{EXP} in Equation (2), with $EXP_{t-1} = sector_{t-1}$ as exposure variable. Disposable income is defined as the sum of labor income, transfers and capital income, after tax. Expenditure is imputed from the household budget constraint as explained in the main text. All variables are scaled by disposable income in $t-1$. The classification of sectors is based on Statistics Norway (2017), see Section 2 for details.