

The Saving and Employment Effects of Higher Job Loss Risk*

Ragnar E. Juelsrud[†] and Ella Getz Wold[‡]

October 2020

Abstract

In this paper we use Norwegian tax data and a novel natural experiment to isolate the impact of job loss risk on saving behavior. We find that a one percentage point increase in the separation rate increases liquid savings by roughly 1.3 - 1.7 percent at the individual level. The response is driven by low-tenured workers, who face the largest increase in job loss risk. Further, we show that an increase in savings due to higher job loss risk at the local level is associated with lower employment in non-tradable industries not directly affected by the shock, also after controlling for lower demand from affected firms and labor mobility across sectors. The results are consistent with a household demand channel of recessions.

Key words: Precautionary savings, household finance, recessions

JEL Codes: D14, E20, E21

*This working 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 want to thank Scott Baker, Gauti Eggertsson, John Friedman, Neil Mehrotra, Plamen Nenov, Matthew Turner, Jesse Shapiro and seminar participants at Brown University, BI Norwegian Business School, Norges Bank and OsloMet for valuable comments and suggestions. We also thank an anonymous referee for the Norges Bank Working Paper Series.

[†]Norges Bank: ragnar@juelsrud.no

[‡]Norges Bank. ella.getz.wold@gmail.com

1 Introduction

Saving rates tend to increase during recessions, and the increase following the recent financial crisis was especially large and long-lived. This has sparked a new interest in both the determinants and effects of higher saving rates during periods of economic distress. Policymakers and academics have linked the increase in savings to higher economic uncertainty – often pointing to an increase in *job loss risk*.¹ Higher job loss risk increases the volatility of expected future income, while at the same time lowering the level. Both of these effects may induce people to save more, and hence consume less. The reduction in consumption implies a reduction in household demand, making the saving response a potential amplifier of economic downturns.

A recent theoretical literature emphasizes the importance of higher savings in response to increased job loss risk in amplifying economic downturns (Bayer et al. 2015, Challe and Ragot 2016, Challe et al. 2017, Ravn and Sterk 2016, Ravn and Sterk 2017). However, little is known about the empirical effect of job loss risk on savings during periods of economic distress. Estimating this effect is challenging, as it requires both an exogenous increase in job loss risk and a strategy to isolate the impact of job loss risk from other recession effects, such as falling house prices. Further, evaluating whether the saving response reduces overall employment through the household demand channel requires a strategy to separate the general equilibrium effects of higher household savings from other forces affecting employment.

In this paper we use administrative panel data from Norway and a novel natural experiment to study the impact of higher job loss risk on savings and local employment. The sudden collapse of the international oil price in 2014 led to an exogenous increase in job loss risk for certain occupations and regions. Using our individual level data, we can compare workers who live in the same area, but who are subject to different changes in job loss risk, allowing us to separate the effect of higher job loss risk from other local recession effects. We find that a one percentage point increase in job loss risk increases liquid savings by 1.3 - 1.7 percent. This observed saving response is consistent with a standard consumption-savings model (Huggett, 1996) in which agents have constant relative risk aversions with a relatively high CRRA parameter (above one), and is inconsistent with experimental estimates of the CRRA coefficient, e.g. Holt and Laury (2002). The estimated saving response can explain roughly 80 percent of the observed increase in liquid saving rates during the oil price collapse, indicating that job loss risk is an important determinate of saving rate changes during recessions.

In order to evaluate the effect of higher savings on local employment, we focus on employment in industries not directly affected by the oil price collapse. After accounting for lower demand from directly affected industries and labor mobility across sectors, we document that non-tradable sector employment declines more in regions in which the increase in individual savings is larger – consistent with the household demand channel. Our results suggest that lower household demand

¹See for example Mody et al. (2012) and Pistaferri (2016).

resulting from higher job loss risk can explain about forty percent of the decline in non-tradable sector employment.

The tax data includes information on income and wealth, and can be merged with labor market data as of 2000. We thus have detailed information on labor market status and occupation, which will be important in identifying individual level job loss risk. We use the 2014 oil price collapse to obtain an exogenous increase in job loss risk which differs across occupations. The occupational group with the largest increase in job loss risk is engineers. As engineers have at least 1 - 3 years of higher education, we compare engineers to other high skilled workers in order to obtain a suitable control group. Prior to the oil price collapse, engineers and other high skilled workers have similar levels of job loss risk, averaging roughly one percent per year. Following the oil price collapse, job loss risk for engineers increases sixfold, while job loss risk for other high skilled workers increases only moderately. As a robustness exercise, we also use an alternative control group consisting of high skilled *government* workers, who did not experience any increase in job loss risk.

In order to control for other local recession effects which potentially affect savings, we start by comparing individuals with different changes in job loss risk, but who live in the same area, in a dynamic difference in difference regression. Specifically, we define the oil region to be the two counties in the South-West of Norway which employ an unproportionally high share of oil workers. By comparing engineers and other high skilled workers who live in the oil region, we can control for any local recession effects which are common across these two groups. In order to evaluate the sign and the magnitude of other local recession effects, we compare the baseline results to an alternative specification in which the control group consists of high skilled individuals *not* residing in oil counties.

The results show an annual increase in savings for engineers relative to other high skilled workers of roughly \$1,300, or just above three and a half percent. Scaling this by the increase in job loss risk, we find that a one percentage point increase in the separation rate increases savings by 1.3 - 1.7 percent. Reassuringly, the increase in savings is driven by low-tenured engineers, who experienced the largest increase in job loss risk. We also document that low-tenured engineers have more favorable outcomes than high-tenured engineers *conditional* on job loss, suggesting that the saving response is not caused by potentially confounding human capital depreciation. Looking only at low-tenured individuals, the increase in savings for every one percentage point increase in job loss risk rises to 1.7 - 3.3 percent. Not controlling for local recession effects has a moderate, but positive impact on the results. This suggests that, if anything, not accounting for other recession effects would cause us to overstate the impact of job loss risk on savings.

When investigating the relevance of the household demand channel, we aggregate the outcome variables to the municipality level and categorize municipalities based on their share of oil sector engineers. We restrict the sample to the municipalities in the oil region. Not surprisingly, municipalities with a higher number of affected individuals experience a relative increase in average savings.

In order to evaluate the overall employment impact of higher savings, we consider industries *not* directly affected by the shock.

Identifying the general equilibrium effects of the risk induced increase in savings on employment is challenging, as there are several factors at work. In order to guide our analysis, we therefore construct a simple model to clarify the key factors through which a negative shock to the oil sector affects non-oil employment. The model highlights three channels. First, a negative shock to the oil sector implies lower demand for the firms producing inputs to the oil sector. Second, unemployed oil workers may switch to other sectors, potentially crowding out employment in these sectors. Third, household demand is reduced, as a result of i) increased savings resulting from higher job loss risk and ii) reduced consumption resulting from actual job loss.

We account for lower firm demand by using input output data and network analysis from Acemoglu et al. (2016). Further, we account for labor mobility across sectors by calculating the number of unemployed oil workers who switch to other sectors following the oil price collapse. In the tradable sector, the cross-sectional increase in unemployment is not statistically significant, consistent with the tradable sector being less sensitive to local demand. In the non-tradable sector however, there is a statistically significant increase in unemployment, which is not fully accounted for by lower firm demand or labor mobility across sectors – suggesting that some of the increase in unemployment is due to lower household demand.

While we do not have an identification strategy to separate the impact of lower consumption resulting from *realized* unemployment from lower consumption resulting from higher job loss *risk*, we argue that the latter is quantitatively more important. Back of the envelope calculations suggest that the total consumption loss resulting from the risk channel is more than four times as large as the total consumption loss resulting from realized unemployment. The reason being that, although unemployed individuals have larger consumption declines, there are relatively few of them compared to the many affected workers who keep their jobs but face an increase in risk. As a result, the decomposition exercise suggests that the risk induced increase in savings can explain about forty percent of the increase in non-tradable sector unemployment. We thus conclude that the data is consistent with job loss risk being an important amplifier of economic downturns.

1.1 Literature review

Several papers study the connection between job loss risk and savings. Most of these papers do not focus on economic downturns specifically, and use either subjective unemployment beliefs (Guiso et al. (1992), Carroll and Dunn (1997), Lusardi (1998)) or future unemployment spells (Chetty and Szeidl (2007), Basten et al. (2016), Hendren (2017)) to capture job loss risk. This has the benefit of not confounding the impact of risk with other recession effects, but does not necessarily capture the impact of job loss risk on savings *conditional* on macroeconomic distress. In order to address endogeneity concerns, this literature has often used mass layoffs to control for within-firm

selection into unemployment (see for example Basten et al. (2016)). However, as pointed out by Hilger (2016), this does not control for potential across-firm selection.

In order to obtain an exogenous increase in job loss risk, Fuchs-Schündeln and Schündeln (2005) use the German reunification as a natural experiment. The German reunification implied a permanent and “once-in-a-lifetime” reassignment of job loss risk across occupations however, and is therefore less relevant for understanding the implications of business cycle variations in job loss risk. An alternative approach is to instrument for (changes in) job loss risk with variables such as region of residence, occupation, sector and demographic characteristics. This approach is taken in Carroll et al. (2003) and Harmenberg and Oberg (2016). Due to the many variables used as instruments, it is not clear exactly what is driving the variation in risk. However, given that region and occupation are important determinants, the exercise may be conceptually similar to the one in this paper. We expand upon the analysis in these papers by separating the impact of job loss risk from other local recession effects, such as falling house prices.

Our analysis is also related to papers which use VARs to identify the impact of different types of uncertainty shocks on consumption and output, such as Alexopoulos et al. (2009), Jurado et al. (2015), Fernández-Villaverde et al. (2015), Leduc and Liu (2016), Larsen (2017) and Basu and Bundick (2017). While these papers typically use uncertainty indices constructed from volatility in variables such as stock prices, our uncertainty measure is job loss risk – which has both a variance effect and a level effect on expected future income. Basu and Bundick (2017) show that an uncertainty shock decreases both consumption and output, and develop a model in which output falls due to an increase in desired savings. We complement their analysis, by providing micro-level evidence in favor of this mechanism. Note that the VAR exercise cannot rule out that output falls as a direct response to the shock, and that this reduces employment and hence consumption. We contribute to this literature by directly showing that savings increase in response to higher uncertainty, and that this increase occurs prior to the employment fall. Further, we explicitly account for intersectoral linkages and labor mobility across sectors, and show that the employment fall is found in the non-tradable sector only, supporting the household demand channel.

Finally, our paper relates to a literature which uses cross-sectional variation to uncover evidence on the local household demand channel. Mian and Sufi (2014) show that employment in the non-tradable sector declines in response to a fall in housing net worth, Verner and Gyongyosi (2018) show that employment in non-exporting firms declines in response to an increase in household debt resulting from a sudden currency crisis and Chodorow-Reich et al. (2019) show that non-tradable sector employment increases in response to higher stock market wealth. In addition to studying a job loss risk shock rather than a net wealth shock, we contribute to this literature by considering savings directly and documenting that the saving response precedes the employment decline - thereby offering further support for the household demand channel.

2 Data and institutional background

We use administrative data which covers the universe of Norwegian tax filers. The main outcome variable is liquid savings, measured by bank deposits. However, we also consider other financial assets. The tax data can be merged with labor market data as of 2000, providing us with detailed information on labor market status and occupation. The latter will be important in identifying which individuals experience an increase in job loss risk.

The tax data is a panel data set, covering the period 1993 to 2017. The data is annual, and variables are measured at the end of the year. It contains information on income from different sources, including transfers and taxes. We define individuals as unemployed if they receive unemployment benefits in a given year. In addition to income data, there is also rich information on household wealth. We observe financial wealth in the form of bank deposits and other financial assets. Real wealth is reported as primary housing wealth, secondary housing wealth and other real wealth. Prior to 2010 the value of real wealth which is reported for tax purposes is substantially below market value. From 2010 and onward, efforts are made to correctly report the market value of housing wealth. The data set also contains information on total debt, allowing us to measure net wealth.

Our main outcome variable is liquid savings, measured by bank deposits. As bank deposits is a highly liquid and safe financial asset, it seems like a good candidate for precautionary saving. However, we will also consider any adjustments that come through other financial assets or real wealth. Bank deposits are reported by the bank, and include saving accounts, checking accounts, fixed term deposits etc. Bank deposits do not include investments in bonds and direct and indirect holdings of stocks, which belong to *other financial assets*. Close to 100 percent of the sample have some positive holdings of bank deposits in a given year, while a substantially lower share own other financial assets or real wealth.

Income is reported and taxed individually in Norway, whereas wealth is reported individually and taxed at the household level. Our unit of analysis is the individual, and so we cannot rule out that there is some misreporting of wealth within the household. However, we expect bank deposits to be relatively well measured also at the individual level, as it is reported by the bank and must be reported as belonging to the owner of the bank account. We follow much of the existing literature in focusing exclusively on men (see for example Basten et al. (2016)).

The tax data can be merged with labor market data as of 2000. Our full data set therefore covers the period 2000 to 2017. From the labor market data we obtain detailed information on occupation and sector, which is important for our identification strategy. The matched firm-worker data also allows us to calculate the observed tenure for each worker, which will be useful for identifying groups with especially large increases in job loss risk.

Occupation is only observed for employed individuals, and there are some instances of employed individuals not having a reported occupation. We therefore define an individual as belonging to an

occupation o if we observe the individual as being employed in that occupation for at least one of the three years leading up to the shock. Similarly, the unemployment rate for a sector o is defined as the unemployment rate for individuals in that occupation. We use the same type of assignment rule for assigning workers to a sector, and for calculating sector level unemployment rates.²

We divide employed individuals into three occupational groups. The first group consists of engineers and civil engineers. The former requires 1-3 years of higher education, whereas the latter requires a minimum of four years of higher education. The second group consists of individuals who are employed in occupations requiring some higher education, and who are not engineers. We refer to this group as *other high skilled* workers. Managers, professionals, technicians and associate professionals belong to this group. In total, close to 50 percent of employed individuals are categorized as being either engineers or other high skilled workers, see Table 11 in Appendix B. The remaining working individuals are employed in occupations which do not require higher education, and are referred to as low skilled.

In addition to using only men, we make some further sample restrictions. First, we use a 25 percent random sample of the tax filing population. Second, we exclude individuals with business income in order to obtain a well defined concept of job loss risk. Third, we only include individuals who are employed at baseline and who can be matched to an occupation in one of the three years leading up to the shock. We also winsorize the variables at the 99 percent level, following Basten et al. (2016).

Summary statistics for the three occupational groups are reported in Table 1. Nearly everyone owns some bank deposits, although the average and median holdings are substantially larger for high skilled workers than for low skilled workers. Engineers and other high skilled workers hold similar amounts. Among the high skilled, just above 60 percent own other financial assets, and other high skilled workers own somewhat more of these assets than engineers. As there is a substantial share of managers in this group, this could perhaps reflect that some of their labor compensation takes the form of financial assets. Among the low skilled, less than 40 percent own other financial assets. Also note that these other financial assets appear relatively skewed within groups, with average holdings far exceeding median holdings.

Engineers and other high skilled workers also look similar in terms of real wealth. Exactly 76 percent in both groups are homeowners, compared to less than 50 percent for low skilled workers. Just above 70 percent in both groups have positive net wealth. The average wage income among engineers is roughly \$95,000, which is somewhat higher than for other high skilled workers, and substantially higher than for low skilled workers. High skilled workers are older than low skilled workers, but engineers and other high skilled workers have similar average and median ages at 44 to 45 years. We thus conclude that engineers and other high skilled workers look fairly similar along

²In the appendix, we also show results using an alternative definition of sector level unemployment rates, in which the unemployment rate in sector i is based on the share of unemployed individuals who in their last year of employment were employed in sector i .

observable characteristics, and that both groups have substantially higher wealth and income levels than low skilled workers. For this reason, we restrict the analysis to a comparison of engineers and other high skilled workers.

	Average			Median		
	Engineers	High Skilled	Low Skilled	Engineers	High Skilled	Low Skilled
Bank Deposits	35,900	34,700	19,600	14,200	11,500	5,600
Other Financial Assets	23,800	43,000	11,300	1,600	1,600	0
Prim. Housing Wealth	233,100	252,000	134,100	227,500	238,500	0
Other Real Wealth	44,600	52,300	23,200	8,300	7,700	100
Debt	183,600	197,400	104,200	153,200	161,000	33,200
Wage Income	94,600	85,600	55,400	90,300	78,800	55,600
Age	44	45	38	44	45	37
Bank Deposits > 0 (%)	99	99	98			
Other Fin. Assets > 0 (%)	61	64	39			
Housing Wealth > 0 (%)	76	76	48			
Net Wealth > 0 (%)	72	71	67			
Observations	21,901	74,113	160,223			

Table 1: Summary statistics 2013 in 2015 USD (rounded to closest 100 with USD/NOK 7.5).

2.1 Institutional background

The impact of job loss risk on savings is likely to depend on the unemployment insurance (UI) scheme. That is, not only job loss risk matters, but also the expected income fall upon job loss – or what we might think of as effective job loss risk. OECD data on 2015 replacement rates from the *Tax and Benefit Systems: OECD Indicators* shows that out of the 40 countries included, Norway is ranked as number 18, i.e. close to the OECD median. For comparison, the US is ranked as number 37. All else equal, we would therefore expect job loss risk to have a smaller impact on savings in Norway than in the US.

Norwegian workers who become unemployed are generally entitled to unemployment insurance of 62 percent of pre-unemployment wages for a duration of two years. While there is a requirement to qualify, this is relatively low, and workers with a non-trivial position throughout the calendar year would all be expected to qualify. There is however an upper limit on pre-unemployment wages, meaning that income above a year-specific threshold does not enter into UI calculations. High income earners therefore have an effective replacement rate of less than 62 percent. This turns out to be relevant for our sample, as the treatment group will consist of relatively high-income individuals. Using the year specific thresholds, we calculate an effective replacement ratio of close to 50 percent for our sample.

With regards to the level of job loss risk, Norwegian unemployment rates are among the lowest in the OECD group. Figure 12 in Appendix A depicts harmonized OECD unemployment rates by

country, with the Norwegian unemployment rate typically falling below four percent. While the unemployment rate in Norway has generally been below that in the US, this has changed in recent years. At the same time as the US unemployment rate has recovered from the Great Recession, the oil price collapse in 2014 led to a deterioration of Norwegian labor market conditions. As a result, the unemployment rates in the two countries have been similar for the past three to four years.

When interpreting the results of this study in a broader context, it is useful to keep in mind that the setting is one of relatively low baseline job loss risk, and relatively generous unemployment insurance – although the effective replacement ratio for our treatment group is lower than the national average.

3 Theoretical predictions

When studying uncertainty shocks, macroeconomic models often consider mean preserving spreads to expected future income. While these shocks are compelling in that they isolate the impact of uncertainty, large and salient uncertainty shocks at the household level often have both a variance effect and a level effect. For example, shocks to job loss risk or health risk will typically be “negative” uncertainty shocks, in that they both *increase* the variance of expected future income and *reduce* the level. Both the level effect and the variance effect might contribute to higher savings. Given consumption smoothing, a reduction in expected future income will likely increase savings today. In addition, the increase in the variance of expected future income will contribute to higher savings if there is prudence in the utility function (Kimball, 1990) or if there are potentially binding borrowing constraints. Moreover, it can also lead to increased labor supply.

Under which conditions will an increase in savings resulting from higher job loss risk reduce output? In standard neoclassical models, the increase in savings leads to an increase in investment. In addition, since higher job loss risk induces a precautionary labor supply response, the overall impact on output is positive. Higher job loss risk therefore increases both savings and output, and there is no amplification of economic downturns.

In New Keynesian models with nominal rigidities, the co-movement between savings and output can break down. If prices and interest rates do not fall sufficiently, the increase in investment will be insufficient to make up for the decline in consumption. Further, if labor supply is inelastic, the precautionary labor supply response is eliminated. As a result, higher job loss risk can increase savings, while reducing output. Recently, a handful of papers have studied uncertainty in the form of job loss risk using search and match models with nominal frictions, see Bayer et al. (2015), Challe and Ragot (2016), Challe et al. (2017), Ravn and Sterk (2016) and Ravn and Sterk (2017). In these models, a shock to the separation rate increases job loss risk and induces individuals to save more. *Conditional* on nominal frictions, the increase in savings contributes to a further decline in output.

4 The effect of job loss risk on savings

The first goal of the empirical exercise is to identify the impact of job loss risk on savings. We focus on liquid savings measured by bank deposits. In Figure 23 in Appendix A we document that illiquid assets do not respond significantly to the increase in job loss risk, and that households therefore appear to adjust primarily by changing their liquid asset positions.³⁴ In order to obtain an exogenous increase in job loss risk, we use the 2014 oil price collapse as a novel natural experiment. By comparing liquid savings for individuals with different levels of job loss risk, but who are subject to the same local recession effects, we aim to isolate the impact of job loss risk from other recession effects.

4.1 Natural experiment: The oil price collapse of 2014

The sudden collapse of the oil price in the summer of 2014 led to an exogenous increase in job loss risk for certain regions and occupations. Job loss risk increased mainly in oil producing regions in the South-West of Norway, while the hardest hit occupational group was engineers.

The price of Brent crude oil fell from roughly \$110 to less than \$50 per barrel in the second half of 2014, as seen in Figure 13 in Appendix A. Popular explanations include a slowdown in global demand, especially from China, as well as high supply of shale oil from the US. Tokic (2015) notes that in contrast to the oil price busts of 1991 and 2008, the 2014 bust was not preceded by an oil price spike, and as such was “completely unexpected”. To the best of our knowledge, there has been no suggestions that the oil price collapse of 2014 was in any way related to the Norwegian oil sector, which stands for only about two percent of world production. We thus feel comfortable assuming that the oil price shock was both unexpected and exogenous to the Norwegian economy.

At the start of 2014, the petroleum sector accounted for roughly 25 percent of Norwegian GDP and 40 percent of Norwegian exports. The large and unexpected decrease in oil prices therefore had an adverse effect on the Norwegian labor market. However, as documented below, the negative impact was to a large degree contained to certain regions and occupations.

Regional and occupational variation Oil production is concentrated in the South-West of Norway, as seen from Figure 14 in the appendix. Two out of nineteen counties employ a disproportionately high share of oil sector workers, and we define these two counties as the “oil region”.⁵ The combined population of these two counties in 2014 was close to one million, or 19 percent of the total population.

³The increase in liquid savings relative to illiquid savings is consistent with for instance Bayer et al. (2015), where higher income uncertainty increases preferences for liquidity.

⁴Figure 24 documents that labor income slowly declines following the unemployment risk shock, indicating that precautionary labor supply is unlikely to be an important response in practice.

⁵The two oil counties are *Hordaland* and *Rogaland*, and the largest city in the area is *Stavanger* - sometimes referred to as the oil capital.

The left panel of Figure 1 depicts the percentage point change in unemployment rates by county. The red squares capture the average of the two counties defined as the oil region, while the blue dots capture the remaining seventeen counties. In 2015, the unemployment rate in the oil region increased by more than two percentage points, making it the largest increase in county level unemployment over the past fifteen years. At the same time, most other counties experienced moderate or no increase in unemployment. The unemployment increase in the oil region dampened somewhat in 2016, and started to reverse in 2017. As documented below, the unemployment *level* in the oil region remained elevated in 2017.

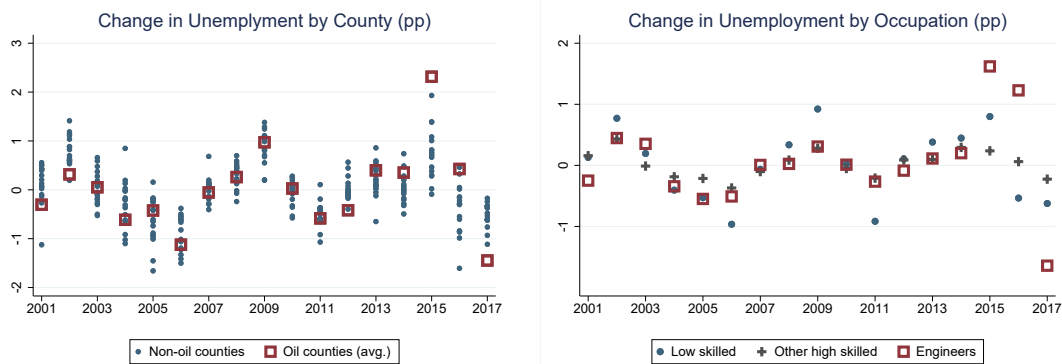


Figure 1: Changes in unemployment rates (pp) by county (left panel) and occupation (right panel).

No other occupational group received as much media attention as engineers following the oil price collapse⁶, and the data suggests that this was indeed warranted.⁷ The tax data contains detailed information on occupations for employed individuals. We categorize individuals as *engineers* if they were employed as engineers in the time leading up to the oil price collapse, i.e. if they were employed as engineers in at least one of the years 2011-2013. The individuals in this group are either civil engineers - which in Scandinavia is a protected title - or engineers. The former requires at least four years of higher education, while the latter requires 1-3 years of higher education. Individuals who do not belong to this group, but who are employed in other occupations requiring higher education, are labeled *other high skilled*.

The right panel of Figure 1 depicts the change in unemployment by occupational group. The change in unemployment rates for low skilled workers is captured by the blue dots. Note that the labor market outcomes of this group seem to be especially cyclical, with high peaks and low busts

⁶Some examples of newspaper headlines: “Statoil is laying off more engineers” *Aftenposten* April 2015, “One out of three engineers are worried about losing their job” *Aftenposten* May 2015, “Union leader for the engineers: Worried unemployment will rise further” *Aftenposten* May 2015, “Solberg [the prime minister] wants to help unemployed engineers” *DN* September 2015. “New report on the oil engineers: Unemployment increased by 342 percent in one year - but many are finding new employment” *E24* March 2016.

⁷The Norwegian Labour and Welfare Administration (NAV) reports unemployment rates for fifteen different occupations, one of which is Engineers & IT workers. According to their data, the increase in unemployment for this group in 2015 was the largest observed increase for any occupational group since their sample starts in 2003.

compared to other workers. The change in unemployment rates for engineers is captured by the red squares, while the change in unemployment rates for other high skilled workers is captured by the plus-signs. These two groups look fairly similar prior to the oil price collapse, but have very different employment outcomes in the year following the shock. In 2015, the unemployment rate for engineers increased by more than 1.5 percentage points - the highest increase observed - while the unemployment rate for other high skilled workers remained roughly unchanged. A similar increase was observed in 2016, with a partial reversal following in 2017. As will become evident in the upcoming analysis, this does not only reflect the geographical distribution of engineers and other high skilled workers.

Salience Figure 1 documented that the oil region experienced a sharp increase in relative unemployment in 2015. Google search data allows us to confirm that not only was the shock quantitatively large, it also appears to have been salient. Search volumes are indexed relative to the maximum search volume in the sample, which is assigned a value of 100. Further, search volumes are measured relative to the total amount of searches in a given area, allowing for meaningful comparisons across geographic areas of different sizes.

The left panel of Figure 2 depicts the volume of searches which Google classifies as belonging to the search category *Brent Blend*, i.e. oil price related searches. The solid red line depicts the volume of oil price related searches in the oil region over time. After the oil price started falling in August 2014, there is an immediate and sustained spike in oil price related searches. As seen from the dashed blue line, the rest of the country follows a very different pattern. Although there is some increase also in other counties, the magnitude is modest compared to that in the oil region. We thus conclude that individuals residing in oil producing areas are especially aware of, and are paying attention to, the collapse in the oil price.

Even though individuals living in affected areas are paying attention to the sudden oil price bust, they need not be aware of the negative consequences for the local labor market. In order to evaluate how salient the shock is in terms of labor market risk, the right panel of Figure 2 depicts the volume of searches which Google classifies as belonging to the search category *Layoff*. Again, we see a rather striking pattern. While there is virtually no increase in layoff related searches in other counties, there is a large and persistent increase in the two oil counties. As before, the increase starts as the oil price begins falling in mid-2014, and then peaks in early 2016. Note that this means that individuals are googling layoffs even before unemployment rates start to rise in the data.⁸

⁸Unemployment rates rise in 2015 according to the tax data, whereas layoff related Google searches increase also prior to 2015. Prior to the oil price collapse in August 2014, the search volume index has an average value of 12. After the oil price collapse, but prior to January 2015, the search volume index has an average value of 28. From January 2015 to December 2017 the search volume index has an average value of 45.

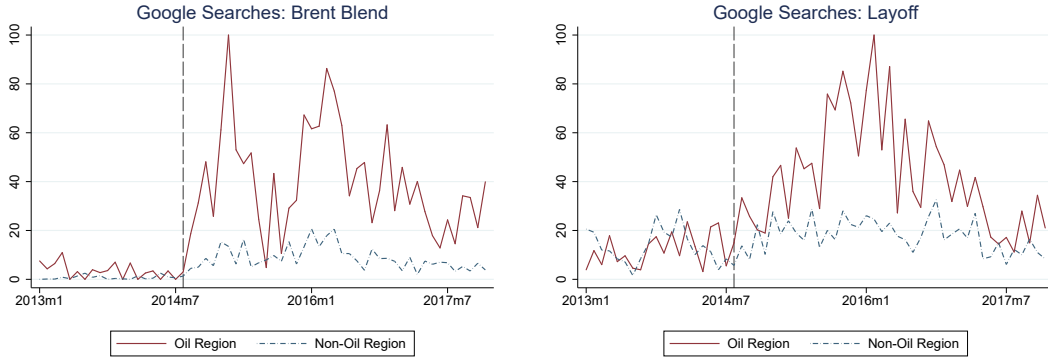


Figure 2: Google search data for the oil region and other counties. The index is set to 100 for the maximum search volume in the sample.

Interestingly, search volumes for layoffs peak in January 2016 (and search volumes for the oil price reaches its second highest value), which is exactly when the oil price reaches its minimum value of \$30 per barrel. Based on the Google search data, we thus conclude that not only are individuals living in oil producing areas immediately aware of the dramatic fall in the oil price, they also seem to understand that this implies an increase in job loss risk.

4.2 Methodology

In order to isolate the impact of job loss risk from other recession effects, we use a difference in difference approach to compare liquid savings for engineers to that of other high skilled workers in the oil region. This within-region comparison allows us to control for the potential impact of other local recession effects on savings, provided that our treatment and control group have similar loadings on the local recession effects. We provide supportive evidence for this in Subsection 4.4. Further, by contrasting the baseline findings to the results from an across-region comparison, we can explicitly evaluate the importance of other local recession effects.

The dynamic difference in difference regression is outlined in equation (1). The main outcome variable Y_{it} is bank deposits for individual i in year t . T_i is an indicator variable equal to one if individual i is in the treatment group, and equal to zero if individual i is in the control group. In the baseline analysis, $T_i = 1$ for engineers residing in oil producing regions, and $T_i = 0$ for other high skilled workers residing in oil producing regions. Treatment status is defined based on the years prior to the oil price collapse. Year fixed effects δ_k are included to capture time-varying aggregate effects which are common to all individuals, while individual fixed effects α_i are included to capture individual, time-constant factors. The coefficients of interest are the β_k 's, which capture the impact of the interaction term between treatment status and year indicator variables. Given that $\beta_k = 0$ for $k < 2014$, the dynamic treatment effect is captured by the β_k 's for $k \geq 2014$. We also estimate the more restrictive difference in difference regression given by equation (2), to obtain

the average treatment effect, in which $I_t^{post} = 1$ if $t \geq 2014$. Standard errors are clustered at the individual level.

$$Y_{it} = \alpha_i + \sum_k \delta_k \mathbf{1}_{t=k} + \sum_k \beta_k (T_i \times \mathbf{1}_{t=k}) + \epsilon_{it} \quad (1)$$

$$Y_{it} = \alpha_i + \sum_k \delta_k \mathbf{1}_{t=k} + \beta_k (T_i \times I_t^{post}) + \epsilon_{it} \quad (2)$$

Because we are interested in the impact of job loss risk, rather than the impact of realized unemployment, we restrict the baseline analysis to only include individuals who are not (yet) unemployed.⁹ However, we also show results using the full sample to avoid potential selection issues, and the estimated responses are similar.

In order to evaluate the importance of local recession effects in determining savings, we complement the baseline analysis with an across-region specification. That is, we compare engineers in oil producing regions to high skilled workers residing outside of oil producing regions. The results from this comparison should reflect both the impact of higher job loss risk and the impact of other local recession effects, such as a relative decline in house prices. Contrasting these results with the baseline findings allows us to also evaluate the sign and magnitude of the impact of other recession effects on savings.

Selection into unemployment Before presenting the results, we briefly discuss the issue of selection into unemployment. In a typical event study in which job loss risk is identified by future unemployment, an important concern is that there is an individual level shock which is causing the upcoming job loss and affecting current saving behavior. This concern is strongly mitigated in our setting, as job loss is caused by an exogenous fall in the oil price – and not by an individual level shock. However, that does not mean that job loss (risk) is randomly distributed within the affected groups. For instance, as we show in the upcoming analysis, engineers with low tenure are more likely to experience job loss than engineers with high tenure. Our estimated saving response will reflect the behavior of people who experience a relatively large increase in job loss risk, which is not necessarily representative of the total population.

We show in Appendix C that after controlling for tenure, other observable characteristics are not informative in predicting which engineers experience job loss following the oil price collapse. Further, we show that a simple model based on observable characteristics has substantially less power in explaining job loss following the oil price collapse than in “normal” times. Hence, to the extent that observable characteristics are relevant for evaluating selection into unemployment, there appears to be relatively less selection following the oil price collapse. This suggests that studying

⁹Specifically, we condition on job loss not occurring between 2014 and 2017, and show saving responses up until 2016. As a result, our sample only consists of individuals who will not become unemployed for at least another year, and are therefore unlikely to have received severance pay or any extraordinary income related to job loss.

the effects of unemployment during times of crisis may get us closer to identifying population representative responses, and hence improving external validity.

4.3 Results

The empirical results confirm that higher job loss risk increases liquid savings. Reassuringly, the increase in savings is driven by low-tenured workers, who experienced an especially large increase in job loss risk. Not accounting for local recession effects produces larger estimates, suggesting that other recession effects might also contribute to higher savings.

Figure 3 depicts the unemployment rate and the separation rate in the oil region over the period 2001-2017, for engineers and other high skilled workers. We include both the unemployment rate and the separation rate, as they capture different aspects of unemployment risk. The separation rate is defined as the probability of transitioning from employed to unemployed. While the separation rate captures the risk of job loss, the unemployment rate is closer to capturing the total risk of unemployment – as it also reflects the job finding rate. As seen from the figure, engineers and other high skilled workers have very similar unemployment and separation rates prior to 2014. This is important as it alleviates the concern that individuals are selecting into our control and treatment groups based on differences in risk aversion, a selection issue studied in detail in Fuchs-Schündeln and Schündeln (2005).

The unemployment rate for engineers increases from an average of roughly one percent prior to the oil price collapse, to a peak of almost seven percent after the oil price collapse. There is some increase in unemployment rates also for other high skilled workers. However, the increase is moderate compared to engineers. In the robustness section, we use an alternative control group consisting only of high skilled *government* workers. This group experienced virtually no increase in job loss risk following the oil price collapse. Reassuringly, the results from this exercise are similar, suggesting that spillovers to the control group is not a concern.

The separation rate is depicted in the right panel of Figure 3. As was the case for the unemployment rate, the separation rate for engineers and other high skilled workers is similar prior to 2014. Post-2014, there is a large and sustained increase in the separation rate for engineers relative to that of other high skilled workers. Note that the separation rate increases by a similar magnitude as the unemployment rate in 2015, but by a smaller amount in 2016. This suggests that the initial increase in unemployment is driven almost exclusively by the separation rate, while a decline in the job finding rate is important in explaining the subsequent increase. By 2017, the separation rate for engineers has almost fallen back to its pre-crisis level, whereas the unemployment rate remains more visibly elevated.

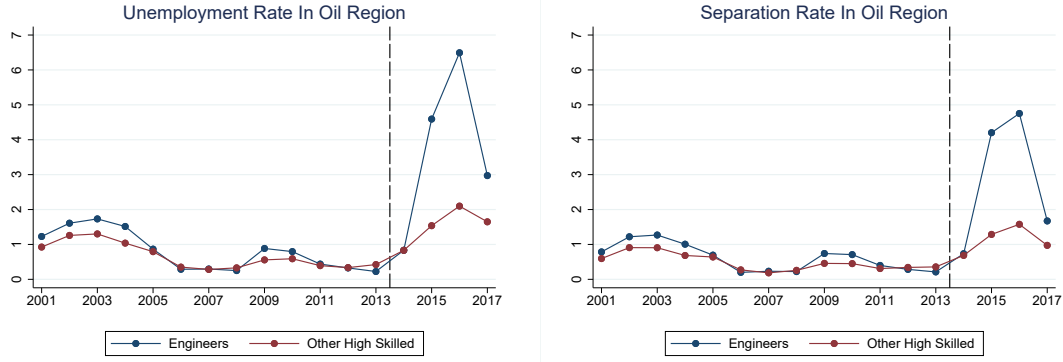


Figure 3: Unemployment rate and separation rate (%) for engineers in the oil region and other high skilled workers in the oil region.

The left panel of Figure 4 depicts bank deposits for engineers and other high skilled workers over time. Bank deposits for the two groups follow each other closely up until 2013, at which time there is a divergence which persists until 2016. Reassuringly, the divergence appears to be driven by an above trend increase in bank deposits for engineers rather than a below trend increase in bank deposits for other high skilled workers. Regression results from estimating equation (1) with $Y_{it} = \text{Bank Deposits}_{it}$ are depicted in the right panel of Figure 4. The baseline sample consisting only of job keepers is captured by the blue dots, while the alternative sample in which we do not condition on job status is captured by the red dots. As seen from the graph, the two samples result in very similar estimates. The pre-2014 coefficients are all very close to zero in magnitude and not statistically significant, suggesting that the parallel trend assumption is satisfied prior to the oil price collapse. In 2014, the coefficient is positive at roughly \$1,300 and statistically significant, implying that engineers in the oil region increased their bank deposits relative to that of other high skilled workers in the oil region.

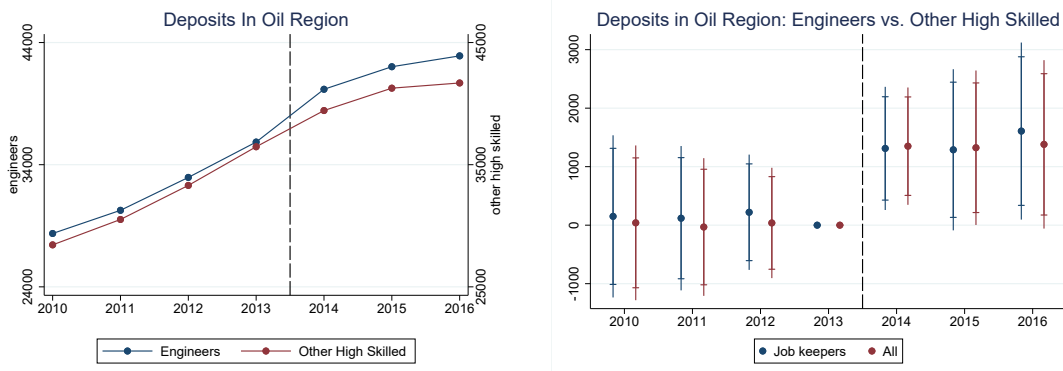


Figure 4: Bank deposits for engineers in the oil region relative to other high skilled workers in the oil region for job keepers. Right panel: coefficient estimates from estimating equation (1) i) using only job keepers and ii) not conditioning on job status.

In Appendix A Figure 15, we further show that while the average saving response occurs in 2014, engineers who lose their job in 2016-2017 increase savings mainly in 2015.

The results in Figure 4 are further summarized in Table 2. As seen from the first column, engineers increased their bank deposits by roughly \$1,300 or 3.6 percent in 2014. In order to scale the saving response, we estimate the increase in unemployment rates and separation rates using a simple difference in difference regression as the one outlined in equation (2). Following search and matching models such as Ravn and Sterk (2017), we use the next period increase in uncertainty to scale the current period saving response. Scaling the saving response by the relative increase in the unemployment rate, we find that a one percentage point increase in the unemployment rate increases liquid savings by 1.3 percent. Alternatively, we can scale the increase in bank deposits by the change in the separation rate, which similarly suggests that a one percentage point increase in the job loss rate increases liquid savings by 1.3 percent.

Results averaging over 2014-2016 are reported in the second column of Table 2, and show a similar increase. Focusing on the 2014 results has the advantage of capturing the initial saving response, which occurred before unemployment started to increase in the data and before any policy changes were implemented or even discussed. This makes it less likely that other forces are behind the relative increase in savings for engineers. However, the shock increased both in size and salience over time, and so we also include results which reflect the saving response in the following years. This has very little impact on the level increase in savings, but slightly increases the scaled responses.

	(1)	(2)
	Bank Deposits	Bank Deposits
$T_i^{2013} \times I_t^{post}$	1,279** (2.26)	1,285** (2.24)
Increase in Bank Deposits (%)	3.63	3.65
per pp increase in unemployment rate (%)	1.27	1.34
per pp increase in separation rate (%)	1.30	1.71
Sample period	2010-2014	2010-2016
Clusters	19,042	18,450
N	93,714	126,954

t statistics in parentheses. Std. errors clustered at individual level. Regressions include individual and year fixed effects.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Bank deposits. Within oil region analysis. Regression results from estimating equation (2) for job keepers.

Tenure While engineers residing in oil regions experienced a general increase in job loss risk after 2013, the increase in risk was not uniformly distributed. In particular, individuals with low tenure faced an especially large increase in the probability of job loss. The Basic Agreement between

the Norwegian Confederation of Trade Unions (LO) and the Confederation of Norwegian Business and Industry (NHO) clearly states that tenure should be an important factor in deciding who gets laid off as a result of cutbacks or restructuring (§ 8-2 *Seniority in the event of dismissal due to cutbacks*). The seniority or tenure principle should only be departed from when “there is due reason for this”. Given that low-tenured individuals faced a particularly large and salient increase in job loss risk, one would expect these individuals to have larger saving responses.

We estimate tenure by calculating the number of years an individual has worked at the same firm. Because the individual tax data can only be matched to employer information as of 2000, the maximum observed tenure prior to the oil price collapse is fourteen years. In 2013, the median observed tenure of engineers residing in oil regions is six years. We thus define individuals with less than six years tenure in 2013 as having *low tenure*. Figure 16 in Appendix A confirms that tenure is indeed an important predictor of unemployment. While the unemployment rate for high-tenured engineers increases to a maximum of almost four percent, the unemployment rate for low-tenured engineers increases to a maximum of nearly ten percent. A similar difference is seen in separation rates.

The results are reported in Table 3, and show that the saving increase is driven by low-tenured workers. Low-tenured engineers increase their liquid savings by roughly \$2,200, while the increase for high-tenured engineers is not statistically significant. As low-tenured engineers have lower holdings of bank deposits to begin with, the percentage increase exceeds seven percent. Scaling the saving response by the relative increase in the unemployment rate, we find that a one percentage point increase in the unemployment rate increases liquid savings by 1.71 percent. Alternatively, a one percentage point increase in the job loss rate increases liquid savings by 1.74 percent. The relative saving response is higher when averaging over the 2014-2016 period, reaching an increase of 3.3 percent for every one percentage point increase in the separation rate.

Relative to the increase in job loss risk, the saving response of low-tenured engineers is higher than the baseline results. This is consistent with the simulation results in Engen and Gruber (2001), in which the percentage effect of risk on savings increases in the level of risk.

	(1)	(2)
	Bank Deposits	Bank Deposits
$T_i^{2013} \times I_t^{post}$	414.1 (0.46)	122.0 (0.14)
$T_i^{2013} \times Tenure_i^{low} \times I_t^{post}$	2,235** (2.00)	2,953*** (2.63)
Increase in Bank Deposits (%) (low tenure)	7.81	10.3
per pp increase in unemployment rate (%)	1.71	2.51
per pp increase in separation rate (%)	1.74	3.29
Sample period	2010-2014	2010-2016
Clusters	18,725	18,134
N	92,141	124,787

t statistics in parentheses. Std. errors clustered at individual level. Regressions include individual and year fixed effects.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Bank deposits by tenure. Within oil region analysis. Regression results from estimating equation (2) for job keepers by tenure.

Other recession effects Local economic downturns can affect saving behavior not only through increased job loss risk. For instance, falling house prices may induce people to cut back on consumption and increase savings. One could also imagine a local recession leading to negative sentiments or beliefs, which might make individuals save more regardless of their employment prospects. In the baseline analysis we did a within region comparison, in order to control for such local recession effects. In this section we explore different specifications in order to evaluate whether these other recession effects are quantitatively important in terms of affecting saving behavior.

The first column in Table 4 simply reproduces the baseline results, in which engineers in the oil region are compared to other high skilled workers in the oil region. In the second column, we compare engineers in the oil region to other high skilled workers everywhere. Finally, in the third column we compare engineers in the oil region to high skilled workers in the non-oil region. Both the coefficient estimates and the scaled rise in liquid savings increase as we move to the right in the table. This suggests that other recession effects are, if anything, contributing to higher saving rates, and that not accounting for these effects would lead us to overstate the impact of job loss risk on savings. However, the difference between the coefficient estimates is not statistically significant.

Note that the quantitative importance of local recession effects is likely to vary, and we do not attempt to measure the size of such effects for our given shock. It is therefore possible that other local recession effects would have larger implications for saving behavior in a different setting, simply because the other local recession effects would themselves be larger.

	(1)	(2)	(3)
	Bank Deposits	Bank Deposits	Bank Deposits
$T_i \times I_t^{post}$	1,279** (2.26)	1,494*** (3.13)	1,542*** (3.20)
Increase in Bank Deposits (%)	3.63	4.24	4.38
per pp increase in unemployment rate (%)	1.27	1.33	1.34
per pp increase in separation rate (%)	1.30	1.34	1.35
Control group: high skilled workers...	in oil region	in all regions	in non-oil region
Sample period	2010-2014	2010-2014	2010-2014
Clusters	19,042	76,798	63,885
N	93,714	379,457	315,702

t statistics in parentheses. Std. errors clustered at the individual level. Regressions include individual and year fixed effects.

* $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Bank deposits. Across region analysis. Regression results from estimating equation (2) for job keepers.

Interpreting the increase in liquid savings Bank deposits are a safe and highly liquid form of savings, and therefore a good candidate for precautionary saving. Basten et al. (2016) find that individuals respond to future unemployment by increasing both the level and the share of safe assets in their portfolio. We have rerun the baseline analysis using total financial wealth as the dependent variable, and the results are reported in Table 12 in the appendix. The increase in total financial wealth is virtually the same as the increase in bank deposits, indicating that non-deposit financial wealth was kept roughly unchanged. There was also no statistically significant decline in housing wealth or other real wealth for engineers relative to other high skilled workers following the oil price collapse.

Because there is no decrease in other forms of wealth – and no relative increase in wages – we find it likely that the increase in liquid savings implied a reduction in consumption. While we cannot rule out that there were other adjustments which we do not observe, we find the 2014 increase in savings especially convincing. At this point there was still no increase in actual unemployment, and the full extent of the oil price collapse was not yet known. As a result, there were no policy measures being discussed at this time. We therefore find it highly probable that the increase in liquid savings implied a reduction in consumption.

4.4 Robustness

In this section, we start by discussing an alternative driver of the estimated saving increase - human capital depreciation - and argue that this interpretation is not supported by data. We proceed by showing that our results are robust to two alternative specifications. First, we change the treatment group to only consist of engineers who work in the oil sector, as these individuals may have been particularly effected by higher job loss risk. Second, we change the control group to only consist of

high skilled government workers, who did not experience any increase in job loss risk following the oil price collapse. We further show that the estimated saving response is unlikely to be driven by wealth effects or selection into occupation based on risk aversion.

Human capital depreciation The sudden oil price collapse and the resulting macroeconomic consequences may have altered oil workers perception of their future earnings potential, and induced them to increase current saving in order to smooth consumption. We refer to this alternative explanation as “human capital depreciation”. While we find it plausible that oil workers did indeed adjust their expectations of future earnings, we argue that the data is inconsistent with this channel driving our results.

We have documented that the saving response is driven by low-tenured engineers, who had the largest increase in job loss risk. In order for the human capital depreciation channel to be driving our saving estimates, low-tenured workers must also have larger losses of human capital than high-tenured workers. Given that low-tenured workers are younger and likely to be more mobile both in terms of geography and industries, this would be somewhat surprising. It would also be inconsistent with results in Couch and Placzek (2010), which show that older workers with greater employment tenure in Connecticut experienced annual earnings reductions five years after job loss more than double those of younger workers.

Table 5 reports outcomes for low tenure and high tenure workers conditional on job loss, and show that low-tenured workers outperform high-tenured workers. Conditional on job loss, low-tenured engineers in 2017 have an income equal to 72 % of their 2013-income, compared to 68 % for high-tenured workers. The difference is larger when considering only wage income. Among low-tenured engineers, 66 % are employed as wage takers in 2017, compared to 57 % for high-tenured workers. This is not explained by high-tenured workers transitioning into retirement, as a higher share of high-tenured workers are still unemployed in 2017. One reason why low-tenured engineers do better might be their willingness to move in order to gain employment. This is supported by 89 % of low-tenured engineers still living in the oil region in 2017, compared to 96 % of high-tenured engineers. Hence, the data does not support the hypothesis that low-tenured engineers have larger losses of human capital - if anything, they outperform their high-tenured peers after job loss. We thus argue that human capital depreciation is unlikely to be the main driver of the estimated saving increase.

	Low tenure	High tenure
Income 2017 / Income 2013	72 %	68 %
Wage income 2017 / Wage income 2013	56 %	44 %
Employment 2017	66 %	57 %
Unemployment 2017	15 %	21 %
Oil region residence 2017	89 %	96 %

Table 5: Engineers in the oil region who lost their job between 2014 and 2016.

Engineers in the oil sector So far, our classification of individuals into treatment and control groups have relied only on occupations. However, we also know in which sector individuals work. We now change the treatment group to only contain engineers which were employed in the oil sector prior to 2014. This leads to, if anything, a higher saving response than in our baseline results.

Statistics Norway defines the oil sector to contain what they refer to as petroleum sectors and petroleum related sectors. The petroleum sector includes the following sectors: extraction of crude petroleum and natural gas (06), support activities for petroleum and natural gas extraction (09.1), transport via pipeline (49.5) and support activities pipeline (52.215). In addition, Statistics Norway defines petroleum related sectors to include the following industries: building of oil-platforms and modules (31.113), installation and completion work on platforms and modules (30.116) and offshore supply terminals (52.223). According to Statistics Norway, around 84,000 individuals were employed in the oil sector in 2014 (Ekeland, 2017) – which constitutes just above three percent of all employed workers. However, a high number of individuals work in industries which produce output used in the oil sector, but which are not included in this definition. Attempts by Statistics Norway to calculate the number of workers directly or indirectly employed in the oil sector based on input output data produces a number of 239,000 – which constitutes just above nine percent of all employed workers (Prestmo et al., 2015). Hence, only 35 % of oil related workers are actually employed in the oil sector.

We follow the standard Statistics Norway definition and create an alternative treatment group, consisting of engineers employed in the oil sector. The new treatment group is thus a subset of our baseline treatment group, while the control group is left unchanged. The time series for unemployment and separation rates for the two groups are depicted in Figure 18 in the appendix, while the evolution of bank deposits is depicted below in Figure 5.

As seen from the left panel of Figure 5, engineers in the oil sector and other high skilled workers have almost identical holdings of bank deposits in the four years leading up to the oil price collapse. Following the oil price collapse, engineers in the oil sector increase their bank deposits relative to other high skilled workers. As reported in Table 13 in the appendix, the increase in bank deposits is somewhat higher than in the baseline - both in absolute value and when scaling the response with the relative increase in job loss risk. A one percentage point increase in the separation rate is now found to increase liquid assets by 1.3-3.4 percent – compared to 1.3-1.7 in the baseline.

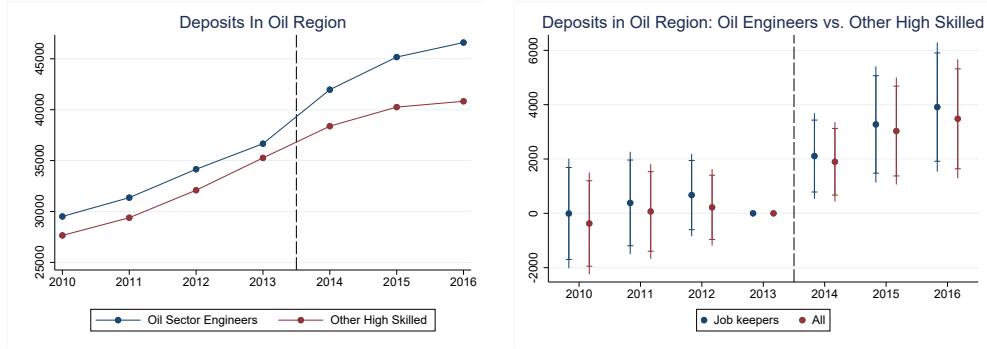


Figure 5: Bank deposits for oil sector engineers in the oil region relative to other high skilled workers in the oil region for job keepers. Right panel: coefficient estimates from estimating equation (1) i) using only job keepers and ii) not conditioning on job status.

Spillovers to the control group The baseline analysis compared engineers residing in oil regions to other high-skilled workers residing in oil regions. It is likely that also the latter group experienced some increase in job loss risk following the oil price shock. Figure 3 showed that although other high-skilled workers in oil regions experienced a very modest increase in unemployment relative to engineers, they too were subject to an increase in job loss risk. This could be because some workers in this group are directly employed in the oil sector and/or because there are spillover effects to other sectors. Note that the largest spillover effects occur for low skilled workers, as alluded to by Figure 1. Hence, this issue is less of a concern when using only high-skilled workers in the control group.

If the impact of job loss risk on saving behavior is homogeneous and linear, spillover effects should not be an issue. To see this note that we are not assuming that there is no increase in job loss risk for the control group. Rather, we are using the difference in job loss risk between the two groups, to scale the impact on liquid savings. If the control and treatment groups have the same underlying linear saving response to a given increase in job loss risk, spillover effects should not affect our estimates. However, if the saving response is non-linear and/or non-homogeneous, spillover effects could be an issue.

To reduce the likelihood that spillover effects are influencing our results we redo the baseline analysis with a control group consisting only of high skilled government workers. This has the benefit of only including individuals whose employment security should not be affected by (short-term) economic conditions, but has the disadvantage of producing a control group with less similar employment outcomes pre-2014. Figure 6 depicts unemployment rates for engineers and high skilled government workers in oil regions. High skilled government workers have virtually no increase in unemployment rates or job loss rates following the oil price collapse, implying limited scope for spillover effects.

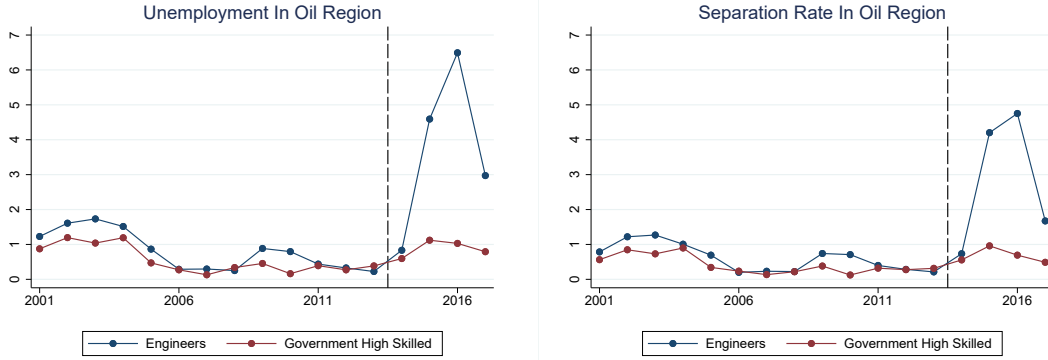


Figure 6: Unemployment rate and separation rate (%) for engineers in the oil region and high skilled government workers in the oil region.

Regression results when using only high skilled government workers in the control group are reported in Table 6. The coefficient estimates are almost unchanged, but the increase in uncertainty is somewhat larger. As a result, a one percentage point increase in the separation rate is found to increase liquid savings by 1.1 percent - compared to 1.3 percent in the baseline. For the 2014-2016 results, a one percentage point increase in the separation rate is found to increase liquid savings by 1.5 percent - compared to 1.7 percent in the baseline. Hence, we conclude that our results are robust to controlling for spillovers to the control group.

	(1) Bank Deposits	(2) Bank Deposits
$T_i \times I_t^{post}$	1,167 (1.60)	1,401* (1.85)
Increase in Bank Deposits (%)	3.31	3.98
per pp increase in unemployment rate (%)	1.06	1.18
per pp increase in separation rate (%)	1.09	1.51
Sample period	2010-2014	2010-2016
Clusters	8,871	8,524
N	43,430	58,338

t statistics in parentheses. Std. errors clustered at individual level. Regressions include individual and year fixed effects.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Bank deposits. Within oil region analysis. Regression results from estimating equation (2) for job keepers using only high skilled government workers in the control group.

House prices Although the within region analysis controls for other “local” recession effects, the definition of local can be disputed. There might still be price differences within the two counties defined as the oil region. For example, engineers and their high skilled peers may live in systematically different areas, thereby being exposed to different changes in house prices. To explore this, we

use house price data on the municipality level from Statistics Norway. This data is not available for the smallest municipalities, but still covers 96 percent of engineers and other high skilled workers residing in the oil region.

Figure 19 in Appendix A depicts average house prices in the oil region over time for engineers and their high skilled peers separately. The change in house prices for engineers and other high skilled workers appears very similar. Prices are roughly constant from 2013 to 2015 for both groups, while house prices in the rest of the country are increasing. House prices in the oil region fall noticeably in 2016, but the decrease is not significantly different across engineers and other high skilled workers.

We also note that the home ownership rates are identical across engineers and other high skilled workers, as showed in the summary statistics in Table 1. Hence, we find it unlikely that house price changes are driving the increase in savings of engineers relative to other high skilled workers, within the oil region.

Other wealth effects If local stock prices are affected, there could also be negative wealth effects coming from financial assets. While there was certainly a decline in stock prices for many oil firms, the overall impact on the Norwegian stock market was limited. As illustrated in Figure 20 in Appendix A, there was some decline in the Oslo Stock Exchange overall index in the second half of 2014, but at an annual level – the relevant level for our tax data – stock prices increased from 2014 to 2015. Moreover, the increase was similar to that of the S&P 500 index in the US. There was a modest fall in stock prices in the following year, but this was also a low growth year for US stock markets. One reason why the oil price collapse appears to have had a relatively modest impact on average stock prices might be the large exchange rate movements, which increased the international competitiveness of Norwegian firms.

Figure 21 in Appendix A shows that there is no decline in the value of other financial assets for engineers or other high skilled workers following the oil price collapse. Note that other financial assets contain not only Norwegian stocks, which might have fallen slightly in value in 2016, but also other assets such as bonds and international stocks - typically held through global mutual funds. If the latter is not hedged against exchange rate movements, the value of these assets would have increased after the oil price collapse.

As long as any wealth effects are constant across the control and treatment group, they cannot be the driving force behind the estimated saving response. As previously discussed, there was no significant change in financial assets for engineers relative to other high skilled workers following the oil price collapse - see Table 12 in the appendix. We thus find it unlikely that the observed increase in bank deposits for engineers relative to other high skilled workers is driven by a negative wealth effect.

Selection into occupations We have used pre-2014 occupations in order to identify groups with different changes in job loss risk. However, occupations are not randomly assigned and engineers

may be systematically different from their high skilled peers. Fuchs-Schündeln and Schündeln (2005) argue that individuals self-select into occupations based on their level of risk aversion, thereby potentially biasing occupation based estimates of precautionary saving. We believe this concern to be of limited importance in our case for two reasons. First, we are comparing two groups which had very similar levels of job loss risk prior to the oil price collapse. As shown in Figure 3, engineers and other high skilled workers had almost identical unemployment rates in the thirteen years leading up to the oil price collapse. Second, we are not simply comparing wealth levels across occupations. Rather, we are considering a sudden change in job loss risk, and the following change in liquid savings. Still, if engineers are less risk averse than the general population, this would mean that the estimated saving response is a lower bound for the population wide response all else equal.

4.5 Implications for aggregate savings rates

A natural question is how the implied increase in savings based on our estimates compares to the observed increase in savings during the oil price collapse. By comparing the two, we can provide a rough back-of-the-envelope calculation of how important the job loss risk effect on savings is for explaining observed increases in saving rates during recessions. In order to do so, we define the liquid saving rate for an individual i at time t as

$$s_{i,t} \equiv \frac{\text{Bank deposits}_{i,t} - \text{Bank deposits}_{i,t-1}}{\text{Wage income}_{i,t}} \quad (3)$$

The change in the saving rate from time $t - 1$ to t is then simply

$$\Delta s_{i,t} = s_{i,t} - s_{i,t-1} \quad (4)$$

Focusing on our sample of (job-keeping) oil-engineers, the average observed change in saving rates from 2013 to 2014 is

$$\overline{\Delta s_{2014}} = 1.51$$

While our estimated saving increase of roughly \$1,300 scaled by the average 2014 wage corresponds to a saving rate increase of

$$\overline{\Delta s_{2014}}^{\text{estimated}} = \frac{1,300}{\overline{\text{Wage income}_{2014}}} = \frac{1,300}{105,700} = 1.23$$

Hence, these simple calculations indicate that the job loss risk effect on savings can explain $1.23 / 1.51 = 81$ percent of the observed increase in saving rates. We thus conclude that the increase in job loss risk is a quantitatively important driver of higher savings.

4.6 Implications for structural parameters

How does the estimated saving increase line up with the saving response implied by a workhorse consumption-saving model? In order to investigate this, we consider a continuous time adoption of Huggett (1996) following Achdou et al. (2017) and focus on the average partial equilibrium saving response to an increase in job loss risk. The analysis is outlined in Appendix D. Here, we simply convey the main conclusions.

In general, the saving response is especially sensitive to at least two household parameters: the discount factor ρ and - given constant relative risk aversion - the CRRA-coefficient γ . We assume that households have CRRA utility and consider a wide range of possible combinations of γ and ρ . We then ask which combinations of these parameters, if any, can generate a saving response consistent with our empirical estimate.

Considering a 95 percent confidence interval around our baseline result, we find that a relatively wide range of parameters are consistent with our findings, see Figure 30 in Appendix D. An important take-away is however, that in order to match our empirical results, the risk aversion parameter γ has to be relatively large and above one. Specifically, for reasonable values of the households discount rate, a γ of at least 1 - 4 is required to generate a saving response of the right magnitude. This implied value of γ is in line with often-used values in macroeconomic models, but substantially higher than experimental estimates of the CRRA coefficient, see e.g. Holt and Laury (2002).

5 The effect of higher savings on local employment

In this section we investigate whether the risk induced increase in savings identified in the previous section may have led to a decrease in employment at the municipality level. This is the second mechanism needed to produce amplification, as discussed in Section 3. Identifying general equilibrium effects such as this is challenging, as there are several effects at play. We start by developing a simple model, highlighting the different channels through which employment is likely to be affected. We proceed by attempting to quantify these channels, arguing that household demand is an important driver of higher unemployment in the *non-tradable* sector.

While the oil price collapse led to a substantial increase in unemployment in the oil sector, other sectors were also affected. We follow the sector definitions used in Mian and Sufi (2014), and consider the tradable and the non-tradable sector separately. The tradable sector is defined as industries with export shares in the top 20th percentile. The non-tradable sector consists of retail, food services and accommodation, as well as construction and real estate firms.¹⁰ As seen in Figure 7, unemployment in the tradable sector increased by roughly 60 % as much as in the

¹⁰Mian and Sufi (2014) considered the construction sector as a separate sector as they were studying the effect of a housing shock. Because we are studying a different shock, we group the non-tradable sectors into one sector. Industries which are not classified as belonging to the oil sector, the tradable sector or the non-tradable sector are categorized as belonging to the residual sector *other*.

oil sector, whereas unemployment in the non-tradable sector increased by somewhat less. We will not attempt to explain the aggregate increase in unemployment in non-oil sectors however, rather we will focus on *cross-sectional* differences in unemployment rates. This implies that any factor which is constant within the oil region (and within a sector), such as the exchange rate, should be accounted for.

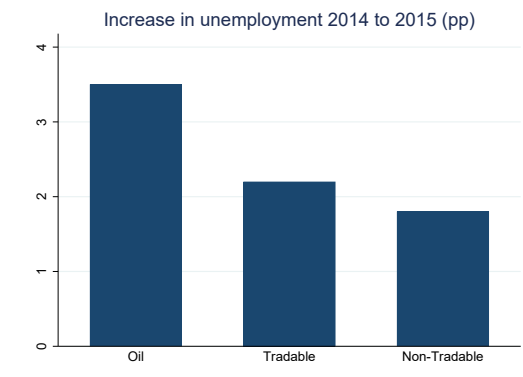


Figure 7: Increase in oil region unemployment by sector from 2014 to 2015 (pp).

In order to obtain cross-sectional variation, we aggregate saving and labor market outcomes to the municipality level and restrict the sample to only include the 59 municipalities in the oil region. To construct our treatment and control groups, we calculate the share of oil sector engineers by municipality at the baseline. Municipalities with an above median share of oil engineers are classified as *oil intensive*, whereas municipalities with a below median share of oil engineers are classified as *non-oil intensive*. As shown in Table 14 in the appendix, bank deposits increase by roughly \$300 per person in oil intensive municipalities relative to non-oil intensive municipalities, confirming that our individual level results hold also at the municipality level.

Figure 8 depicts unemployment rates for oil intensive and non-oil intensive municipalities by sector. While unemployment rates are similar across municipalities up until 2014, there is a divergence in 2015 as unemployment rates increase faster in oil intensive municipalities. This cross-sectional difference is especially clear in the non-tradable sector, generally assumed to be the most dependent on local household demand. Regression results are reported in Table 15 in the appendix. We focus on the 2015 results in order to minimize the amount of confounding factors, but show results for 2015-2016 in the appendix. Relative unemployment in oil intensive municipalities increases by 1.9 percentage points in the non-tradable sector in 2015, while the increase in the tradable sector is substantially smaller and not statistically significant.

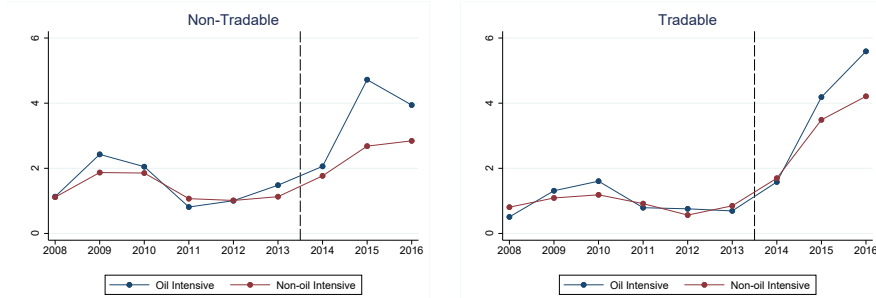


Figure 8: Oil region unemployment by sector and municipality type (%).

Our goal for the rest of this section is to evaluate to what extent the increase in household savings caused by higher job loss risk has contributed to the cross-sectional increase in non-tradable unemployment.

5.1 Model

In order to clarify how a drop in demand from oil workers can increase local unemployment, we develop a simple conceptual framework with different sectors and municipalities. The setup is an extension of Mian and Sufi (2014), and considers additional channels relevant to the shock studied here. Specifically, we consider two channels which the data suggests are relevant: i) lower demand from oil sector firms may reduce employment in other sectors through intersectoral linkages, and ii) (unemployed) oil sector workers may become employed in other sectors, potentially crowding out employment in these sectors.

We outline the model here, and refer to Appendix E for further details. Consider a continuum of municipalities, in which municipalities are indexed $m \in \mathcal{M}$. Each municipality consists of three sectors producing three distinct outputs; non-tradable goods (N), tradable goods (T) and oil goods (O). N and T are consumed by households, while O is sold internationally. Each municipality is populated by two types of workers; a measure one of non-oil workers (indexed no) and a measure ℓ_m of oil-workers (indexed o). To capture cross-sectional variation in exposure to lower demand from oil-workers, we allow ℓ_m to vary across municipalities and assume that labor cannot move between municipalities. Within a municipality, non-oil workers can work in both the T and the N sector. Subject to some frictions specified below, oil workers can switch to other sectors. The local price of non-tradable goods is denoted $P_{m,N}$, the price of tradable goods is denoted P_T , while nominal international demand for oil is given by \bar{y} .

Tradable and non-tradable production The output of the tradable good $y_{m,T}$ and the non-tradable good $y_{m,N}$ in each municipality is produced using labor. Specifically

$$\begin{aligned}
y_{m,T} &= \tilde{e}_{m,T} \\
y_{m,N} &= \tilde{e}_{m,N}
\end{aligned}$$

where $\tilde{e}_{m,i}$ is the effective employment in municipality m and sector i . Effective employment encompasses both non-oil workers and any (displaced) oil workers who have switched sectors. Total effective labor supply outside of the oil sector is given by

$$\tilde{e}_{m,T} + \tilde{e}_{m,N} = E_m \quad (5)$$

Perfect labor mobility implies that wages are equal across the tradable and non-tradable sectors, which further implies that $P_{m,N} = P_N = P_T \forall m$. We assume no input-output linkages between the T and N sector for simplicity, but relax this assumption when taking the model to data.

Non-oil workers Workers consume goods from the tradable and the non-tradable sector. We follow Mian and Sufi (2014) and assume that the total nominal demand of non-oil workers is some quantity given by $\mathcal{D}_{m,no}$. Non-oil workers have Cobb-Douglas preferences with a non-tradable consumption share α , which yields demand functions

$$P_N C_{m,N,no} = \alpha \mathcal{D}_{m,no} \quad (6)$$

$$P_T C_{m,T,no} = (1 - \alpha) \mathcal{D}_{m,no} \quad (7)$$

Oil firms Each municipality has a measure f_m of price-taking oil firms. We assume that oil firms are identical across municipalities, and produce using a Cobb-Douglas production function

$$y_O = X_N^{a_N} X_T^{a_T} e_O^{a_L} \quad (8)$$

with $a_N + a_T + a_L = 1$. Oil goods are thus produced using intermediary inputs from the non-tradable sector X_N and the tradable sector X_T , in addition to labor. Market clearing in the oil sector is given by

$$y_O \leq \bar{y} \quad (9)$$

The combination of a Cobb-Douglas production function and the market clearing condition yields the following conditional factor demands

$$P_T X_T = a_T \bar{y} \quad (10)$$

$$P_N X_N = a_N \bar{y} \quad (11)$$

$$w_O e_O = a_L \bar{y}$$

Oil-workers Oil workers have one unit of labor which they supply inelastically. Their income is therefore $\mathcal{D}_o = a_L \bar{y}$. We assume that also oil workers have Cobb Douglas preferences with a non-tradable consumption share α , which yields demand functions

$$P_N C_{N,o} = \alpha \mathcal{D}_o \quad (12)$$

$$P_T C_{T,o} = (1 - \alpha) \mathcal{D}_o \quad (13)$$

Market clearing in the market for tradable and non-tradable goods The following market clearing conditions pin down the prices P_N and P_T

$$y_{m,N} = C_{m,N,no} + \ell_m C_{N,o} + f_m X_{m,N} \quad \forall m \quad (14)$$

$$\int_{m \in \mathcal{M}} y_{m,T} dm = \int_{m \in \mathcal{M}} C_{m,T,no} dm + \bar{\ell} C_{T,o} + \bar{f} X_T \quad (15)$$

where $\bar{\ell} \equiv \int_{m \in \mathcal{M}} \ell_m dm$ and $\bar{f} \equiv \int_{m \in \mathcal{M}} f_m dm$, i.e the average fractions of oil-workers and oil firms across all municipalities. In Appendix E, we use these market clearing conditions to characterize the baseline equilibrium.

Labor market mobility between sectors While non-oil workers can work in both the tradable and non-tradable sectors, we impose some frictions on the labor mobility of oil workers within municipalities. We want to capture the fact that oil-workers and non-oil workers to some extent are imperfect substitutes in the labor market and, as a result, employment of oil-workers might affect the employment of non-oil workers. Specifically, if employment in the oil sector is reduced, some of that employment reduction might induce oil workers to seek employment in other sectors, potentially reducing employment in these sectors¹¹. We refer to this mechanism as crowding out resulting from labor mobility across sectors.

In order to capture this channel in a reduced-form and parsimonious way, we assume the following labor market clearing condition

¹¹Note that this is sensitive to how we measure sector level unemployment. Because we measure unemployment in sector s based on the individuals who are employed in sector s at baseline, an inflow of workers from other sectors might increase the measured “baseline” unemployment in sector s .

$$\tilde{e}_{m,i} = e_{m,i} + \ell_m \kappa_i (1 - e_O) \quad (16)$$

i.e., lower employment in the oil sector, e_O , increases total effective labor supply in sector i . The strength of this effect is captured by the sector specific labor friction parameter κ_i and the amount of oil workers in the municipality ℓ_m .

Oil sector shock Suppose now that prices are fixed at their baseline values and there is a shock to oil demand $d\bar{y}$. When prices can not adjust, employment in the different sectors is entirely demand driven. As shown in Appendix E, cross-sectional differences in the employment of non-oil workers in the *tradable* sector between two municipalities m and m' is then given by

$$\frac{de_{m,T} - de_{m',T}}{d\bar{y}} = (\ell_m - \ell_{m'}) \kappa_T \frac{\partial e_O}{\partial \bar{y}} \quad (17)$$

i.e. the difference is due to differences in the local labor crowding out effect. In contrast, the cross-sectional difference in *non-tradable* employment is given by

$$\begin{aligned} \frac{de_{m,N} - de_{m',N}}{d\bar{y}} &= \frac{\alpha}{P_N^*} \left(\frac{\partial \mathcal{D}_{m,no}}{\partial \bar{y}} - \frac{\partial \mathcal{D}_{m',no}}{\partial \bar{y}} \right) + (\ell_m - \ell_{m'}) \frac{\alpha}{P_N^*} \frac{\partial \mathcal{D}_o}{\partial \bar{y}} \\ &+ (\ell_m - \ell_{m'}) \kappa_n \frac{\partial e_O}{\partial \bar{y}} + (f_m - f_{m'}) \frac{a_N}{P_N^*} \end{aligned} \quad (18)$$

In the non-tradable sector, there are at least four sources of cross-sectional differences in employment. First, the change in household demand coming from *non-oil* workers can differ across municipalities (the first term in equation (18)). Second, the change in household demand coming from *oil* workers can differ across municipalities (the second term in equation (18)). Third, as in the tradable sector, the crowding out of labor resulting from lower employment in the oil sector can differ across municipalities (the third term in equation (18)). And finally, the effect of lower demand from oil firms might differ across municipalities (the fourth term in equation (18)).

To simplify equation (18) somewhat, assume that cross-sectional differences in demand from non-oil workers is driven by cross-municipal differences in the employment of non-oil workers in the non-tradable sector. If we assume that non-oil worker demand is equal to labor income earned from employment in the non-tradable and the tradable-sector, this is equivalent to assuming that cross-sectional variation in the evolution of employment of non-oil workers in the tradable sector is small (which is confirmed by data, see Figure 8 and Appendix Table 15). That is,

$$\frac{\partial \mathcal{D}_{m,no}}{\partial \bar{y}} - \frac{\partial \mathcal{D}_{m',no}}{\partial \bar{y}} \approx w_N \left(\frac{de_{m,N}}{d\bar{y}} - \frac{de_{m',N}}{d\bar{y}} \right) \quad (19)$$

where $w_N = P_N$ is the non-tradable sector wage. Given this assumption, the differences in

non-tradable employment between two municipalities is given by

$$\frac{de_{m,N} - de_{m',N}}{dy} = \frac{1}{1 - \alpha} \left((\ell_m - \ell_{m'}) \frac{\alpha}{P_N^*} \frac{\partial \mathcal{D}_o}{\partial y} + (\ell_m - \ell_{m'}) \kappa_N \frac{\partial e_O}{\partial y} + (f_m - f_{m'}) \frac{a_N}{P_N^*} \right) \quad (20)$$

When attempting to isolate the impact of lower household demand from oil sector workers on non-tradable sector employment, we will rely on the expression in equation (20).

5.2 Isolating the household demand channel

In order to isolate the impact of lower household demand from oil sector workers on cross-sectional employment in the non-tradable sector, we need to account for at least two other channels. First, we use input output data to account for lower demand from oil sector firms - the last term in equation (20). This is necessary as long as the importance of oil sector firms systematically differs across municipalities in a way that correlates with our treatment and control groups, i.e. $\text{COV}(\ell_m, f_m) > 0$. Second, we directly calculate the importance of labor crowding out by calculating the number of oil workers who switch sectors after becoming unemployed - the second term in equation (20). In addition to these two channels, we further need to consider a multiplier effect which will amplify all the channels. The multiplier arises through the impact of non-tradable sector employment on non-oil workers income, and is captured by the term $\frac{1}{1-\alpha}$ in equation (20).

Given our theoretical setup, any remaining difference in cross-sectional employment in the non-tradable sector will be due to lower household demand coming from oil workers. Note that demand from oil workers can be lower both because employed oil workers save more due to higher risk, and because unemployed oil workers reduce consumption due to lower income. With some back of the envelope conditions we argue that the isolated household demand effect is largely driven by the former channel.

5.2.1 Firm demand

In order to account for lower firm demand, we consider a generalization of our framework following Acemoglu et al. (2016) to identify the sector specific reliance on oil firm demand. Accounting for these intersectoral linkages requires the use of input output data, which we obtain from Statistics Norway.

To accommodate the data, we allow for a more general input-output structure compared to the one used in the conceptual framework. Specifically, we define expenditure shares $a_{ij} \equiv \frac{P_j X_{ij}}{P_j Y_j}$ as the matrix elements in the matrix

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & \ddots & & \\ \vdots & & & \\ a_{n1} & & & a_{nn} \end{bmatrix}$$

The matrix A captures the direct input output links between sectors, with a_{ij} denoting the share of sector j output used by sector i as an input. In our model, the A matrix was particularly simple since only oil firms used inputs from the two other sectors. In order to capture both direct and indirect linkages between sectors, we define the Leontief inverse of this matrix as $H = (I - A)^{-1}$, with matrix elements h_{ij} . To illustrate the importance of also accounting for indirect sectoral linkages, consider two sectors i and j which can be highly connected if sector i uses a substantial share of sector j output as an input. This would be captured by a high a_{ij} . However, the two sectors can also be highly connected if sector k uses a substantial share of sector j output as an input, *and* sector i uses a substantial share of sector k output as an input. This would be captured by a high h_{ij} .

Following Acemoglu et al. (2016) and using the input output matrix, the overall effect of employment in sector i following a shock to oil demand $d\bar{y}$ is

$$d \ln e_i^{\text{firm demand}} = \frac{h_{i,O}}{h_{O,O}} d \ln e_O \quad (21)$$

National input output data The matrix A based on national input output data is reported in Table 7. The bottom row is of special interest, as it tells us what share of production in each sector is used as inputs in the oil sector. For instance, $a_{O,N} = 0.01$ implies that one percent of production in the non-tradable sector is used as an input in the oil sector. This compares to five percent in the tradable sector ($a_{O,T} = 0.05$).

	Non-Tradable	Tradable	Other	Oil
Non-Tradable	0.16	0.05	0.09	0.01
Tradable	0.03	0.19	0.08	0.10
Other	0.14	0.17	0.21	0.04
Oil	0.01	0.05	0.04	0.02

Table 7: Direct sectoral linkages 2013. A matrix.

Total intersectoral linkages are given by matrix H , reported in Table 8. The bottom row now tells us the share of production which is used as inputs in the oil sector – both directly and indirectly. The share of non-tradable production which is used as an input in the oil sector increases from one to three percent when also the indirect linkages are taken into account ($h_{O,N} = 0.03$). In the tradable sector, the share increases from five to eight percent ($h_{O,T} = 0.08$). Because the tradable

sector is the most reliant on demand from the oil sector, it should experience the largest increase in *aggregate* unemployment as a result of lower firm demand.¹² However, in our model, the tradable sector is affected by average rather than local demand, implying that lower firm demand can not account for any *cross-sectional* unemployment increase in the tradable sector. In the appendix, we show results when relaxing this assumption.

	Non-Tradable	Tradable	Other	Oil
Non-Tradable	1.22	0.11	0.15	0.03
Tradable	0.07	1.28	0.14	0.14
Other	0.23	0.30	1.33	0.09
Oil	0.03	0.08	0.06	1.03

Table 8: Direct and indirect sectoral linkages 2013. H matrix.

Adjusting for the regional importance of the oil sector In order to account for lower firm demand, we would ideally want municipality level input output data. Unfortunately, input output data is only available at the national level, and so we adjust the data ourselves to allow for a greater importance of the oil sector in certain areas. Note that if oil intensive and non-oil intensive municipalities in the oil region have the same input output matrices, corporate sector spillovers would not be able to explain any of the cross-sectional increase in unemployment in the non-tradable sector. We therefore allow for the possibility that firms in oil intensive municipalities are more reliant on oil sector demand than firms in non-oil intensive municipalities.

When adjusting the national input output matrix, we assume that connections to the oil sector are proportional to oil sector employment. Because there are 3.5 times as many oil sector employees in oil intensive municipalities as the national average (adjusted for population size), we assume that firms in oil intensive municipalities have 3.5 times as large ties to the oil sector. Similarly, because there are 1.6 times as many oil sector employees in non-oil intensive municipalities, we assume that firms in non-oil intensive municipalities have 1.6 times as large ties to the oil sector. The adjusted input output tables are reported in Tables 16 and 17 in Appendix B.

The employment effects of lower firm demand We now use equation (21) to calculate the predicted increase in non-tradable unemployment resulting from lower oil firm demand. From Figure 7 we know the change in aggregate oil sector unemployment, i.e. $\Delta \ln u_O = -\Delta \ln e_O = 3.5$, where u_O denotes the unemployment rate in the oil sector. Using the matrix elements from the adjusted H-matrix, we can then calculate the predicted increase in unemployment in the non-tradable sector resulting from lower firm demand. The results are reported in Table 9. In the non-tradable sector, lower demand from oil sector firms can explain a relative unemployment increase

¹²It is possible that although the tradable sector produces the most inputs for the oil sector, these are the inputs which oil sector firms are the least likely to cut back on in the short term. In our baseline analysis we follow the model and assume no effect on cross-sectional unemployment in the tradable sector, making this distinction irrelevant.

of 0.19 percentage points.¹³

	Implied unemployment increase (pp)		
	Oil-intensive	Non-oil intensive	Difference
Non-Tradable	$\frac{0.10}{1.11} \times 3.5 = 0.32$	$\frac{0.04}{1.06} \times 3.5 = 0.13$	0.19

Table 9: Predicted unemployment increases from lower firm demand.

From the model and equation (20) we know that the direct impact of lower firm demand will be subject to a multiplier. In the model, the multiplier is determined by the consumption share of non-tradable goods α - which according to the input output data is 34 %. If we instead consider the share of non-tradable output which is consumed by households, we get a similar number of 32 %. Both of these figures result in a multiplier of about 1.5, which is consistent with the broader literature on multipliers. As will be evident from the results however, our findings are qualitatively robust to substantially higher multipliers.¹⁴

Figure 9 illustrates the cross-sectional increase in unemployment rates in 2015.¹⁵ Results based on the period 2015-2016 are similar, and are depicted in Figure 26 in the appendix. Unemployment in the non-tradable sector in oil-intensive municipalities increases by close to two percentages points relative to non-oil intensive municipalities. Lower demand from oil sector firms can explain only about ten percent of this relative increase, or fifteen percent when including the multiplier. In the tradable sector, the relative increase in unemployment is not significant and employment - by assumption - is not affected by local oil firm demand. In Appendix A Figure 27 we show results when relaxing the assumption that local oil firm demand does not affect cross-sectional differences in unemployment in the tradable sector.

¹³In Table 18 we relax the assumption that the tradable sector is not subject to local demand shocks and show the predicted increase in tradable unemployment due to lower firm demand.

¹⁴The relative increase in non-tradable sector unemployment remains statistically significant also when accounting for firm demand and labor mobility across sectors with multipliers up to 4.25 - which is substantially higher than the literature suggests. Ramey (2011) argues that the fiscal multiplier is probably between 0.8 and 1.5, but that “reasonable people can argue that the data does not reject multipliers between 0.5 and 2.0”. More recently, Chodorow-Reich (2019) concludes that his preferred point estimate for a cross-sectional multiplier is 1.8, and suggests a national no-monetary-policy-response multiplier of 1.7 or above.

¹⁵As mentioned in the data section, the sectoral unemployment definition mirrors the individual level definition, in the sense that the unemployment rate in sector i is the unemployment rate among individuals who worked in sector i in at least one of the pre-shock years 2011-2013. In the appendix we show results using an alternative definition, in which the unemployment rate in sector i is based on the share of unemployed individuals who in their last year of employment were employed in sector i , see Figure 25.




Figure 9: Estimated increase in unemployment by sector and predicted unemployment increase from lower firm demand.

5.2.2 Labor mobility across sectors

As illustrated in the above model, some of the increase in unemployment in non-oil sectors could be the result of oil sector workers replacing the existing work force. For example, if an oil sector engineer loses his job, and then gets a new job in construction – *replacing* an existing construction worker – this could increase unemployment in the non-tradable sector (depending on how sector level unemployment rates are defined). In order to determine whether this channel is quantitatively important, we calculate the number of unemployed oil workers who become employed in a non-oil sector in 2015. Note that we implicitly assume that oil sector workers who switch to another sector always *replace* an existing worker, leading to an upper bound on the impact of labor mobility on unemployment.

The amount of oil sector workers who transition to the non-tradable sector in 2015 is equivalent to 0.18 % of total non-tradable sector employment at baseline. These oil sector workers exclusively transition to the construction and real estate sector, with not a single individual transitioning from the oil sector to retail, food services or accommodation. A similar share transition to the tradable sector. The amount of oil sector workers who transition to the tradable sector in 2015 is equivalent to 0.19 % of total tradable sector employment at baseline. Interestingly, if we expand the sample period to also include 2016, the share of oil workers who become employed in the non-tradable sector increases moderately to 0.45 % of total non-tradable employment at baseline, while the share of oil workers who become employed in the tradable sector increases quite dramatically to 3.8 % of total tradable employment at baseline. The 2016-results are illustrated in Figure 26 in the Appendix A.

As with firm demand, note that labor mobility across sector can only explain an increase in cross-sectional unemployment if the amount of oil sector workers switching to other sectors differs across oil intensive and non-oil intensive municipalities. To ensure consistency, we again use oil sector employment to scale the impacts on the two types of municipalities. Given that oil intensive municipalities have 3.5 times as many oil workers as the national average and non-oil intensive

municipalities have 1.6 times as many oil workers, oil intensive municipalities are 2.2 times as oil intensive as non-oil intensive municipalities. Combining this with the fact that in the tradable sector, new employees from the oil sector constitute 0.19 % of total employment, labor mobility across sectors can explain an unemployment increase of 0.14 percentage points in the tradable sector. In the non-tradable sector, the number is 0.13 %, as seen from Table 10.

	Tradable	Non-tradable
Average labor mobility effect (%)	0.19	0.18
Impact in oil intensive municipalities (%)	0.26	0.24
Impact in non-oil intensive municipalities (%)	0.12	0.11
Difference (pp)	0.14	0.13

Table 10: Predicted unemployment increases from labor crowding out in 2015 (pp).

Figure 10 reproduces the previous figure, while adding the effect of labor mobility across sectors. In the non-tradable sector, the combined effect of lower firm demand and labor crowding out sums up to around 25 percent of the total increase in unemployment once the multiplier is added. This implies that, according to our estimates, 75 percent of the relative increase in unemployment in the non-tradable sector can be ascribed to lower household demand in some form.

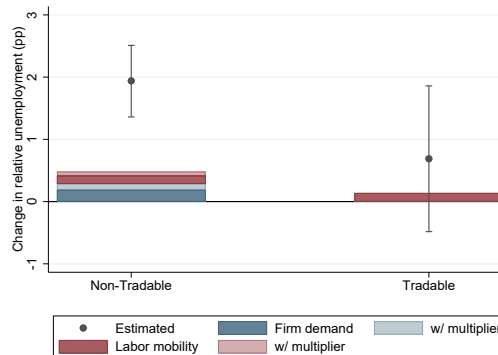


Figure 10: Estimated increase in unemployment by sector, predicted unemployment increase from lower firm demand and labor crowding out.

5.2.3 Household demand

As seen from Figure 10, non-tradable sector unemployment increases by 1.46 percentage points once firm demand and labor mobility across sectors is accounted for. Given equation (20), this increase is attributed to lower demand from oil workers times a multiplier. With the baseline multiplier of 1.5, this implies that lower demand from oil workers increases relative unemployment by 0.97 percentage points - which accounts for fifty percent of the total increase.

Lower demand from oil workers can further be decomposed into two parts. First, oil workers

who are still employed but face higher job loss risk will increase savings and reduce consumption. Second, oil workers who become unemployed will reduce consumption due to lower income.¹⁶ While we cannot isolate the effect of higher realized unemployment from the effect of higher job loss risk, we argue that the latter is quantitatively more important.

To see this, note that for every oil worker who experienced job loss in 2014-2015, twenty-four oil workers kept their job. With some back of the envelope calculations, we can compare the total consumption loss coming from job losers to the total consumption loss coming from job keepers. First, assume that job keepers in the oil sector reduce their consumption by \$1,300 reflecting the results in Table 2 in the previous section. Second, assume that job losers consume all of their after tax income, and that they reduce consumption by 14 percent upon job loss (Browning and Crossley, 2001).¹⁷ This implies that the total consumption loss from the job loss risk channel is $24 \times \frac{1,300}{0.14 \times 51,500} = 4.3$ times as large as the total consumption loss from the realized unemployment channel.

Given that the risk effect is about four times as large as the realized unemployment effect, the former can account for an increase in relative unemployment of 0.78 percentage points. This accounts for forty percent of the total increase in cross-sectional unemployment. Hence, we conclude that, given the assumptions made, not only can lower household demand account for the majority of the cross-sectional increase in non-tradable unemployment, but lower household demand *resulting* from higher job loss risk can explain about forty percent of the increase - making it the largest component of the household demand channel.

6 Conclusion

We have used the oil price collapse of 2014 to identify an exogenous increase in job loss risk for certain segments of the population. By doing a within-region comparison of individuals across different occupations, we estimated that a one percentage point increase in job loss risk increased liquid savings by 1.3 - 1.7 percent. This effect was driven by low-tenured individuals, who faced the largest increase in job loss risk. While low-tenured individuals faced higher job loss risk, their economic outcomes conditional on job loss were relatively more favorable, suggesting that job loss risk rather than human capital depreciation was driving our results. We found no effect on other financial assets, suggesting that the saving response came through bank deposits only.

Further, we showed that unemployment in non-oil sectors increased more in municipalities with

¹⁶If house prices fell in oil intensive municipalities, this could also contribute to lower relative consumption. However, house price growth was roughly zero in 2015, and the difference between oil intensive and non oil intensive municipalities was not statistically significant - see Figure 22 in Appendix A.

¹⁷Several papers use food consumption from the PSID to estimate the consumption drop upon unemployment. These papers generally find consumption falls of less than ten percent, see for instance Chetty and Szeidl (2007). We use a consumption drop of 14 percent as estimated by Browning and Crossley (2001) using Canadian data, as they consider total consumption and Canada and Norway have similar replacement ratios.

more affected individuals and larger saving responses. For the non-tradable sector, the increase in unemployment was not fully accounted for by lower demand from the firm sector, nor by labor mobility across sectors. This suggested that lower demand from the household sector was an important cause. Back of the envelope calculations implied that lower household demand was driven largely by a risk induced increase in savings rather than realized job loss, and that the former channel could account for about forty percent of the total cross-sectional increase in non-tradable sector unemployment.

References

- Daron Acemoglu, Ufuk Akcigit, and William Kerr. Networks and the macroeconomy: An empirical exploration. *NBER Macroeconomics Annual*, 30(1):273–335, 2016.
- Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and wealth distribution in macroeconomics: A continuous-time approach. Technical report, National Bureau of Economic Research, 2017.
- Michelle Alexopoulos, Jon Cohen, et al. Uncertain times, uncertain measures. *University of Toronto Department of Economics Working Paper*, 352, 2009.
- Christoph Basten, Andreas Fagereng, and Kjetil Telle. Saving and portfolio allocation before and after job loss. *Journal of Money, Credit and Banking*, 48(2-3):293–324, 2016.
- Susanto Basu and Brent Bundick. Uncertainty shocks in a model of effective demand. *Econometrica*, 85(3):937–958, 2017.
- Christian Bayer, Ralph Lütticke, Lien Pham-Do, and Volker Tjaden. Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. 2015.
- Martin Browning and Thomas F Crossley. Unemployment insurance benefit levels and consumption changes. *Journal of public Economics*, 80(1):1–23, 2001.
- Christopher D Carroll and Wendy E Dunn. Unemployment expectations, jumping (s, s) triggers, and household balance sheets. *NBER macroeconomics annual*, 12:165–217, 1997.
- Christopher D Carroll, Karen E Dynan, and Spencer D Krane. Unemployment risk and precautionary wealth: Evidence from households’ balance sheets. *Review of Economics and Statistics*, 85(3):586–604, 2003.
- Edouard Challe and Xavier Ragot. Precautionary saving over the business cycle. *The Economic Journal*, 126(590):135–164, 2016.
- Edouard Challe, Julien Matheron, Xavier Ragot, and Juan F Rubio-Ramirez. Precautionary saving and aggregate demand. *Quantitative Economics*, 8(2):435–478, 2017.
- Raj Chetty and Adam Szeidl. Consumption commitments and risk preferences. *The Quarterly Journal of Economics*, 122(2):831–877, 2007.
- Gabriel Chodorow-Reich. Geographic cross-sectional fiscal spending multipliers: What have we learned? *American Economic Journal: Economic Policy*, 11(2):1–34, 2019.

- Gabriel Chodorow-Reich, Plamen T Nenov, and Alp Simsek. Stock market wealth and the real economy: A local labor market approach. Technical report, National Bureau of Economic Research, 2019.
- Kenneth A Couch and Dana W Placzek. Earnings losses of displaced workers revisited. *American Economic Review*, 100(1):572–89, 2010.
- Anders Ekeland. Sysselsatte i petroleumsnæringene og relaterte næringer 2016. 2017.
- Eric M Engen and Jonathan Gruber. Unemployment insurance and precautionary saving. *Journal of Monetary Economics*, 47(3):545–579, 2001.
- Jesús Fernández-Villaverde, Pablo Guerrón-Quintana, Keith Kuester, and Juan Rubio-Ramírez. Fiscal volatility shocks and economic activity. *American Economic Review*, 105(11):3352–84, 2015.
- Nicola Fuchs-Schündeln and Matthias Schündeln. Precautionary savings and self-selection: evidence from the german reunification experiment. *The Quarterly Journal of Economics*, 120(3):1085–1120, 2005.
- Luigi Guiso, Tullio Jappelli, and Daniele Terlizzese. Earnings uncertainty and precautionary saving. *Journal of Monetary Economics*, 30(2):307–337, 1992.
- Karl Harmerberg and Erik Oberg. Consumption dynamics under time-varying unemployment risk. Technical report, Tech. rep., Mimeo, 2016.
- Nathaniel Hendren. Knowledge of future job loss and implications for unemployment insurance. *American Economic Review*, 107(7):1778–1823, 2017.
- Nathaniel G Hilger. Parental job loss and children’s long-term outcomes: Evidence from 7 million fathers’ layoffs. *American Economic Journal: Applied Economics*, 8(3):247–83, 2016.
- Charles A Holt and Susan K Laury. Risk aversion and incentive effects. *American economic review*, 92(5):1644–1655, 2002.
- Mark Huggett. Wealth distribution in life-cycle economies. *Journal of Monetary Economics*, 38(3):469–494, 1996.
- Kyle Jurado, Sydney C Ludvigson, and Serena Ng. Measuring uncertainty. *American Economic Review*, 105(3):1177–1216, 2015.
- Miles S Kimball. Precautionary saving in the small and in the large. *Econometrica: Journal of the Econometric Society*, pages 53–73, 1990.

- Vegard Larsen. Components of uncertainty. 2017.
- Sylvain Leduc and Zheng Liu. Uncertainty shocks are aggregate demand shocks. *Journal of Monetary Economics*, 82:20–35, 2016.
- Annamaria Lusardi. On the importance of the precautionary saving motive. *The American Economic Review*, 88(2):449–453, 1998.
- Atif Mian and Amir Sufi. What explains the 2007–2009 drop in employment? *Econometrica*, 82(6):2197–2223, 2014.
- Ashoka Mody, Franziska Ohnsorge, and Damiano Sandri. Precautionary savings in the great recession. *IMF Economic Review*, 60(1):114–138, 2012.
- Luigi Pistaferri. Why has consumption remained moderate after the great recession. In *paper prepared for Boston Fed Conference.*, 2016.
- J Prestmo, B Strøm, and H Midsen. Ringvirkninger av petroleumsnæringen i norsk økonomi. *SSB Rapport*, 8, 2015.
- Valerie A Ramey. Can government purchases stimulate the economy? *Journal of Economic Literature*, 49(3):673–85, 2011.
- Morten O Ravn and Vincent Sterk. Macroeconomic fluctuations with hank & sam: An analytical approach. 2016.
- Morten O Ravn and Vincent Sterk. Job uncertainty and deep recessions. *Journal of Monetary Economics*, 90:125–141, 2017.
- Damir Tokic. The 2014 oil bust: Causes and consequences. *Energy Policy*, 85:162–169, 2015.
- Emil Verner and Gyozo Gyongyosi. Household debt revaluation and the real economy: Evidence from a foreign currency debt crisis. 2018.

Appendix A: Figures

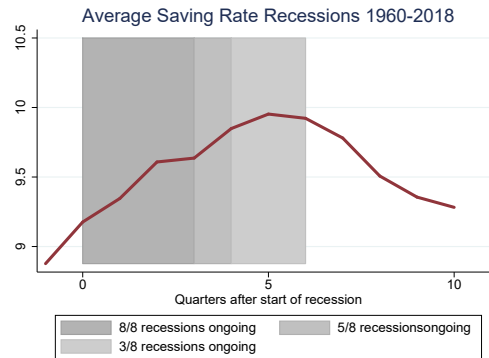


Figure 11: US personal saving rate. Savings as a share of disposable income. Average over past eight recessions (1960-2018). Three quarter moving average. Source: St. Louis FRED database.

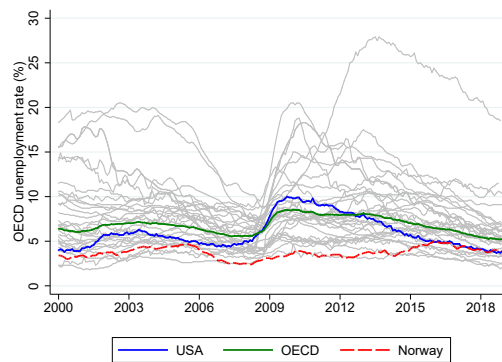


Figure 12: OECD harmonized unemployment rates by country (%).

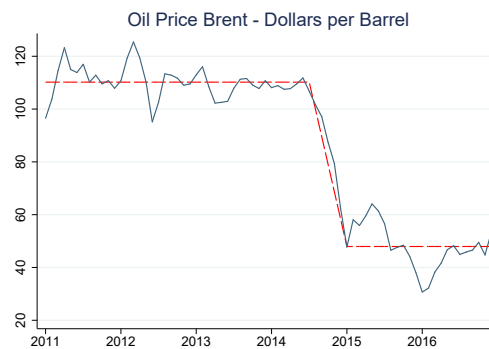


Figure 13: Oil price brent. USD per barrel.




Figure 14: Share of workers employed in the oil sector relative to the share of total workers by county.

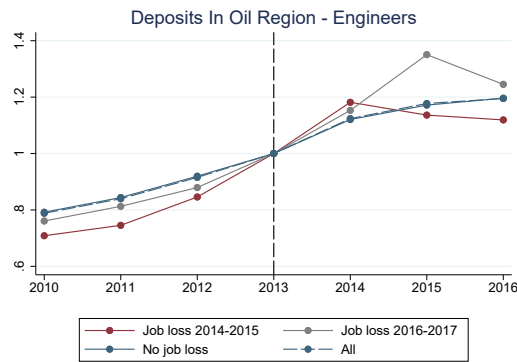


Figure 15: Bank deposits in oil regions for engineers, engineer who did not lose their job following the oil price collapse, and engineers who lost their job in 2016.

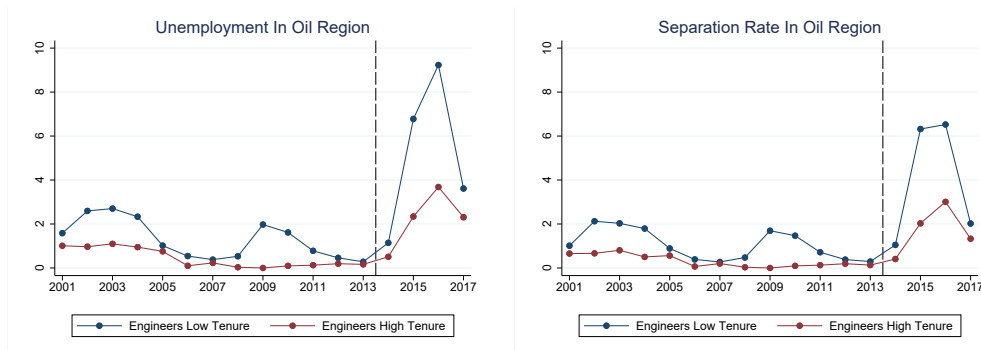


Figure 16: Unemployment rate and separation rate (%) for low tenure engineers in the oil region and high tenure engineers in the oil region.

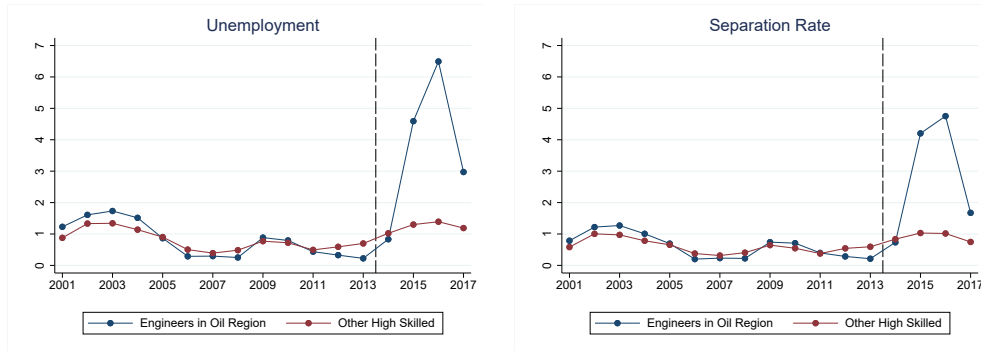


Figure 17: Unemployment rate and separation rate (%) for engineers in the oil region and other high skilled workers in all regions.

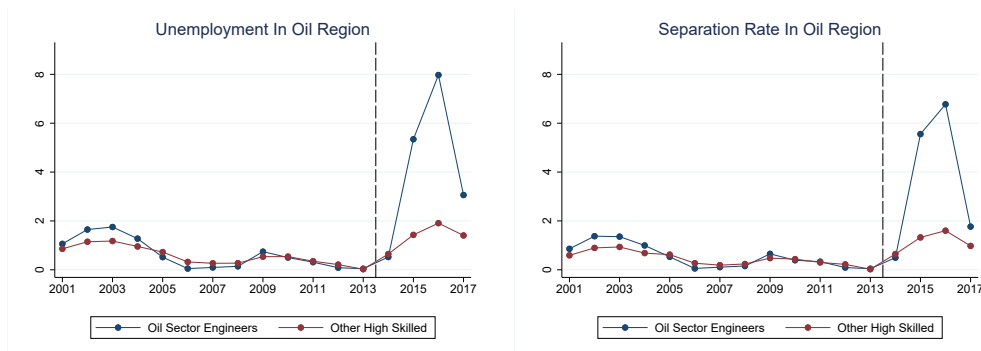


Figure 18: Unemployment rate and separation rate (%) for oil sector engineers in the oil region and other high skilled workers in the oil region.

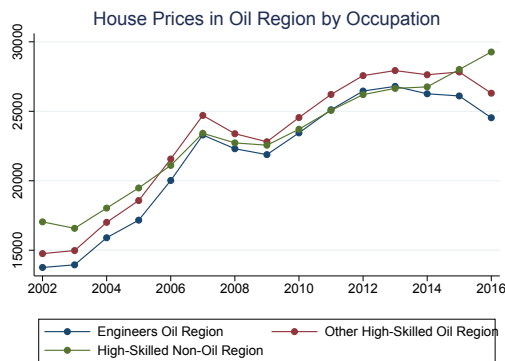


Figure 19: House prices single family homes. Municipality level. Average for engineers and other high skilled workers in the oil region.

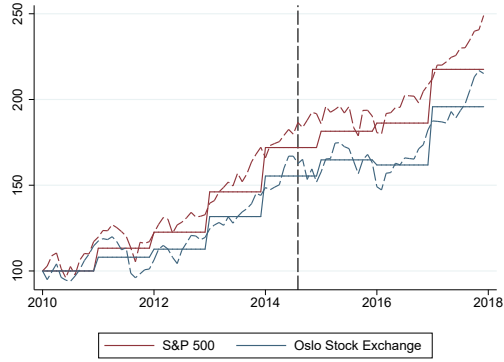


Figure 20: Stock prices. S&P 500 index and Oslo Stock Exchange index. Solid lines are annual data, whereas dashed lines are monthly data.

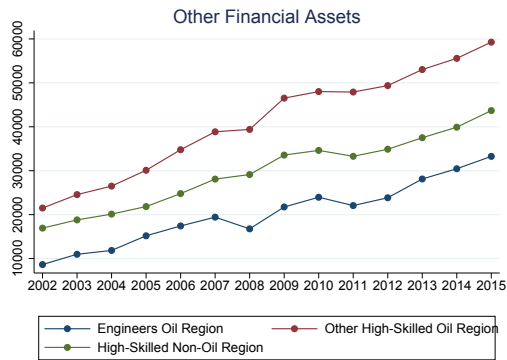


Figure 21: Other financial assets by occupation-region.

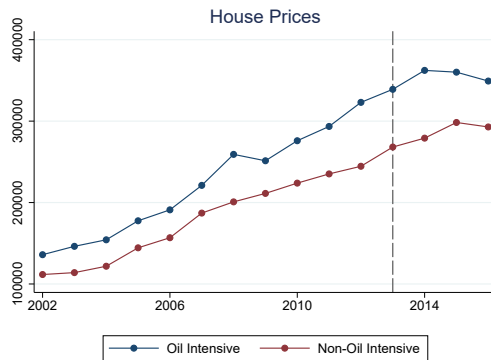


Figure 22: Average house prices for oil intensive and non-oil intensive municipalities.

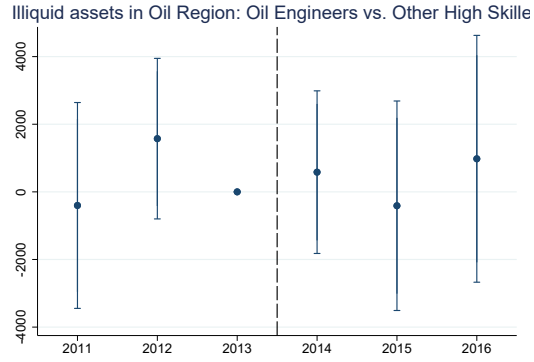


Figure 23: Coefficient estimates from estimating equation (1), using illiquid assets as dependent variable.

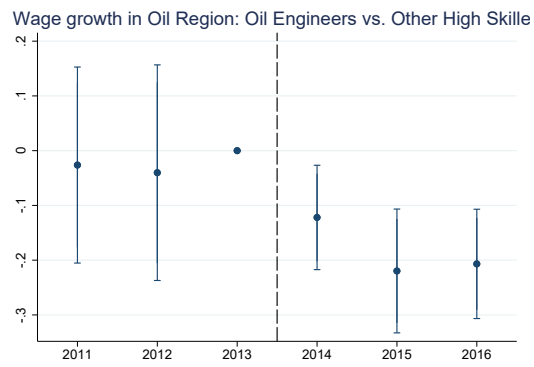


Figure 24: Coefficient estimates from estimating equation (1), using labor income as dependent variable.




Figure 25: Reproduction of Figure (10) using an alternative sectoral unemployment definition, in which unemployment in sector i is based on the share of unemployed individuals who in their last year of employment were employed in sector i .




Figure 26: Reproduction of Figure (10) for 2015-2016.

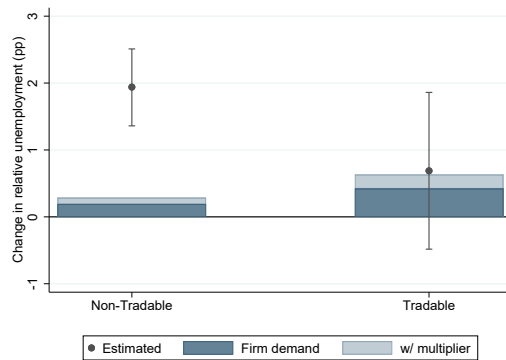


Figure 27: Reproduction of Figure (9) relaxing the assumption that the tradable sector is not affected by lower firm demand.




Figure 28: Reproduction of Figure (10) relaxing the assumption that the tradable sector is not affected by lower firm demand

Appendix B: Tables

Occupations	Education/Skills	Share of Workers (%)
1 - Managers	Not specified	11
2 - Professionals	Min. 4y of higher educ.	15
3 - Technicians/Associate prof.	1y-3y of higher educ.	21
4 - Clerical support workers	High school	6
5 - Service and sales workers	High school	12
6 - Skilled agriculture	High school	1
7 - Craft and related trade workers	High school	17
8 - Plant and machine operators	High school	11
9 - Elementary occupations	Not specified	4
0 - Armed forces and unspecified	Not specified	2

Table 11: Occupations. Occupations 1-3 are classified as high skilled.

	(1)	(2)	(3)	(4)
	Deposits	Deposits	FW	FW
$T_i^{2013} \times I_t^{post}$	1,279**	1,285**	1212	760.7
	(2.26)	(2.24)	(1.25)	(0.71)
Increase in Deposits/FW %	3.63	3.65	1.87	1.17
per pp increase in unemployment rate %	1.27	1.34	0.66	0.43
per pp increase in separation rate %	1.30	1.71	0.67	0.55
Sample period	2010-2014	2010-2016	2010-2014	2010-2016
Clusters	19,042	18,450	19,042	18,450
N	93,714	126,954	93,714	126,954

t statistics in parentheses. Std. errors clustered at individual level. Regressions include individual and year fixed effects.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 12: Bank deposits and total financial wealth (FW). Within oil region analysis. Regression results from estimating equation (2) for job keepers with $Y = \{\text{Bank Deposits, Total Financial Wealth}\}$.

	(1)	(2)
	Bank Deposits	Bank Deposits
$T_i \times I_t^{post}$	1,865**	2,815***
	(2.16)	(3.11)
Increase in Bank Deposits (%)	5.24	7.91
per pp increase in unemployment rate (%)	1.30	1.92
per pp increase in separation rate (%)	1.26	2.38
Sample period	2010-2014	2010-2016
Clusters	14,738	14,303
N	72,656	98,625

t statistics in parentheses. Std. errors clustered at individual level. Regressions include individual and year fixed effects.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 13: Bank deposits. Within oil region analysis Regression results from estimating equation (2) for job keepers, comparing engineers in the oil sector to other high-skilled workers.

	(1)	(2)
	Bank Deposits	Bank Deposits
$T_m \times I_t^{post}$	296.2 (1.27)	335.4* (1.72)
Sample period	2010-2014	2010-2015
N	295	354

t statistics in parentheses. Regressions include individual and year fixed effects.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 14: Bank deposits at the municipality level within the oil region. $T_m = 1$ if municipality m is an oil intensive municipality in the oil region, and $T_m = 0$ if municipality m is a non-oil intensive municipality in the oil region.

	(1)	(2)
	Non-Tradable Unemployment	Tradable Unemployment
$T_m \times I_t^{15}$	1.935*** (6.77)	0.688 (1.18)
Sample period	2010-2015	2010-2015
Clusters	59	59
N	354	349

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 15: Sectoral unemployment rates at the municipality level within the oil region. $T_m = 1$ if municipality m is an oil intensive municipality in the oil region, and $T_m = 0$ if municipality m is a non-oil intensive municipality in the oil region.

	Non-Tradable	Tradable	Other	Oil
Non-Tradable	0.15	0.03	0.07	0.01
Tradable	0.02	0.13	0.06	0.06
Other	0.13	0.11	0.16	0.02
Oil	0.05	0.18	0.13	0.08

Table 16: Direct sectoral linkages 2013. A matrix. Baseline adjustment for oil intensive municipalities in the oil region. Created by taking Table 7 and assuming that $a_{4j}^{adjusted} = 3.5a_{4j} \forall j$ and adjusting all other a_{ij} 's with the same factor $\sum_{i=1}^3 a_{ij}^{adjusted} = x \sum_{i=1}^3 a_{ij} \forall j$ such that the total input share is unchanged $\sum_{i=1}^4 a_{ij}^{adjusted} = \sum_{i=1}^4 a_{ij} \forall j$.

	Non-Tradable	Tradable	Other	Oil
Non-Tradable	0.16	0.04	0.08	0.01
Tradable	0.02	0.18	0.08	0.09
Other	0.14	0.15	0.20	0.04
Oil	0.02	0.08	0.06	0.04

Table 17: Direct sectoral linkages 2013. A matrix. Baseline adjustment for low oil intensive municipalities in the oil region. Created by taking Table 7 and assuming that $a_{4j}^{adjusted} = 1.6a_{4j} \forall j$ and adjusting all other a_{ij} 's with the same factor $\sum_{i=1}^3 a_{ij}^{adjusted} = x \sum_{i=1}^3 a_{ij} \forall j$ such that the total input share is unchanged $\sum_{i=1}^4 a_{ij}^{adjusted} = \sum_{i=1}^4 a_{ij} \forall j$.

	Implied unemployment increase (pp)		
	Oil-intensive	Non-oil intensive	Difference
Non-Tradable	$\frac{0.10}{1.11} \times 3.5 = 0.32$	$\frac{0.04}{1.06} \times 3.5 = 0.13$	0.19
Tradable	$\frac{0.26}{1.11} \times 3.5 = 0.82$	$\frac{0.12}{1.06} \times 3.5 = 0.40$	0.42

Table 18: Predicted unemployment increases from lower firm demand for the non-tradable and tradable sector.

Appendix C: Selection into unemployment

In this appendix, we attempt to quantify the amount of selection into unemployment based on observable characteristics among engineers in the years following the oil price collapse.

We start by evaluating to what extent we can predict job loss during the oil crisis based on baseline characteristics. Specifically, we define an indicator variable $I_i^{jobloss} = 1$ if engineer i experienced job loss in 2015 or 2016, and zero otherwise. We then regress this indicator variable on 2013 characteristics in a probit regression, according to equation (22). Ex-ante, we expect tenure to be an important variable in explaining job loss, as firms are obliged to follow the seniority principle in determining layoffs. Other control variables are captured in X_i , and include age, wage income, total income, housing wealth, real wealth, financial wealth, bank deposits, and debt.

$$I_i^{jobloss} = \alpha + \beta Tenure_i + \gamma X_i + \epsilon_i \quad (22)$$

The regression results are reported in Table 19. As expected, tenure has a negative and significant effect on the probability of job loss. However, after controlling for tenure, information on income, wealth and debt does not have a significant impact on the probability of job loss. The only other variable that is statistically significant – at the ten percent level – is age. When tenure is not included in the regression, both age, financial wealth and debt has a significant effect on the probability of job loss. The pseudo R^2 is low in both cases, but especially so when tenure is excluded from the analysis.

In order to compare the amount of selection during the oil crisis to selection into unemployment

during “normal times”, we repeat the above analysis for job loss prior to the oil price collapse. Specifically, we let $I_i^{jobloss}$ indicate job loss in one of the years 2003-2013 and rerun the regression specified in equation (22). We then compare the pseudo R^2 's to the pseudo R^2 reported in Table 19. The results are depicted in Figure 29. The pseudo R^2 's during the oil crisis is the lowest in the sample, suggesting that the simple statistical model outlined in equation (22) has less explanatory power in predicting job loss during the oil price crisis than in normal times.

Note however, that because we can only calculate tenure back until year 2000, the comparison is somewhat misleading (as the tenure variable contains more information towards the end of the sample). In order to undertake a more fair comparison, we exclude tenure from the model, and redo the analysis. The resulting pseudo R^2 's are depicted in the right panel of Figure 29. The pseudo R^2 during the oil price collapse is now much lower than in normal times, suggesting less selection on observables into unemployment.

	(1) Job Loss	(2) Job Loss
Tenure	-0.0737*** (-10.54)	
Age	0.00402* (1.71)	-0.00490** (-2.22)
Wage Income	0.000000240 (0.22)	-0.000000971 (-0.82)
Total Income	-0.000000957 (-1.09)	-0.000000244 (-0.24)
Primary Housing Wealth	5.77e-08 (0.21)	-6.37e-08 (-0.24)
Real Wealth	-0.000000151 (-0.58)	-0.000000202 (-0.80)
Financial Wealth	-0.000000621 (-1.63)	-0.000000810** (-2.12)
Bank Deposits	9.80e-08 (0.14)	0.000000144 (0.21)
Debt	0.000000202 (1.41)	0.000000236* (1.68)
Constant	-1.082*** (-10.76)	-1.009*** (-10.09)
Pseudo R2	0.0457	0.0133
N	6,732	6,732

t statistics in parentheses
* $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$

Table 19: Regression results from estimating equation (22) with dependent variable $I_i^{jobloss} = 1$ if engineer i experienced job loss in 2015-2016. Probit regression.

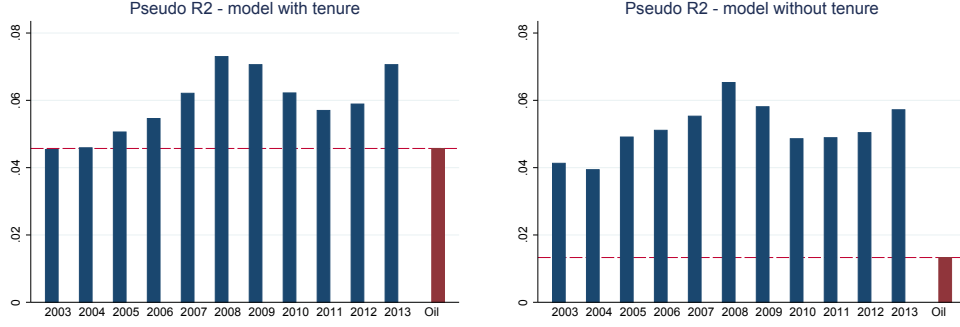


Figure 29: Pseudo R2 from probit regression.

Appendix D: Structural interpretation - details

Model set-up We consider a continuum of individuals with heterogeneous wealth a and income y . Individuals value consumption according to a constant relative risk aversion utility function $u(c_t) \equiv \frac{c_t^{1-\gamma}}{1-\gamma}$ and discount the future at rate ρ . The time-zero present value of future consumption is given by

$$\mathbb{E}_0 \int_0^{\infty} e^{-\rho t} u(c_t) dt \quad (23)$$

Household income y_t evolves stochastically over time according to a two-state Poisson process. Households can either be employed and earn a fixed wage w , or be unemployed. In the unemployed state, labor income is given by $\kappa \times w$, where κ denotes the replacement rate. Transitions from the employed to the unemployed state occur with intensity λ_{EU} and transitions from the unemployed to the employed state occur with intensity λ_{UE} .

Households earn a fixed rate of return r on their asset holdings. Household wealth therefore evolves according to the flow budget constraint

$$\dot{a} = y_t + r a_t - c_t \quad (24)$$

subject to a borrowing limit

$$a_t \geq \underline{a} \quad (25)$$

Subjects maximize lifetime utility (23), subject to equations (24) and (25).

Numerical details We consider the model outlined above in partial equilibrium (i.e. we keep the interest rate fixed exogenously at r). We set the quarterly interest rate $r = 0.015/4$, the wage rate $w = 0.2$ and the replacement rate $\kappa = 0.5$ in line with the data. We then consider two

Description	Parameter value
Wage as employed	$w = 0.2$
Replacement rate	$\kappa = 0.5$
Job finding rate	$\lambda_{UE}^{\text{low risk}} = \lambda_{UE}^{\text{high risk}} = 0.73/4$
Separation rate, low risk economy	$\lambda_{EU}^{\text{low risk}} = 0.0037$
Separation rate, high risk economy	$\lambda_{EU}^{\text{high risk}} = 0.0127$
Borrowing limit	$\bar{a} = -w/3$

Table 20: Structural parameters.

different economies, one “low risk” economy and one “high risk” economy. In both economies, we set the job finding intensity $\lambda_{UE} = 0.73/4$ in line with observed job-finding rates and calibrate the job separation rate λ_{EU} to match the unemployment rate for oil engineers in respectively 2013 (the low risk economy) and 2015 (the high risk economy). This yields $\lambda_{EU}^{\text{low risk}} = 0.0037$ and $\lambda_{EU}^{\text{high risk}} = 0.0127$. Finally, we set the borrowing limit $\bar{a} = -w/3$, i.e. one month of labor income. All of the parameters are summarized in Table 20.

Theoretical saving response Having fixed these parameters, we move on to consider a wide range of values for $\{\gamma, \rho\}$ with the aim of understanding which set of parameters that generates a saving response in line with our findings. Specifically, for both the low risk and the high risk economy, we obtain the steady state distribution of savings for employed individuals, captured by the function $g^i(a)$ and the savings policy function \dot{a}^i for $i \in \{\text{low risk, high risk}\}$. We then compute the saving response to an increase in job loss risk as

$$\Delta s = \int \left(\dot{a}^{\text{high risk}} - \dot{a}^{\text{low risk}} \right) g^{\text{low risk}}(a) da \quad (26)$$

i.e. as the difference in average savings between the high risk and the low risk economy evaluated at the steady-state low risk distribution. For comparability with our empirical estimates, we scale it with the average asset holdings in the low risk economy and define the scaled model-generated saving response as

$$\beta^{\text{model}} \equiv \frac{\Delta s}{\int a g^{\text{low risk}}(a) da} \quad (27)$$

We then compare the model generated saving increase β^{model} with our estimated saving response $\hat{\beta}$. We are interested in combinations of ρ and γ which generate a model-implied saving response β^{model} within the 95 % confidence interval of $\hat{\beta}$. Due to numerical issues when γ is high, we focus on the lower bound of γ that generates a saving response within the 95 % confidence interval of our estimate and how it varies with ρ . We define any combination of $\{\gamma, \rho\}$ which generates a saving response within our 95 % confidence interval as being “data-consistent.”

The set of data consistent parameters $\{\gamma, \rho\}$ is shown in Figure 30. As ρ increases, the lower

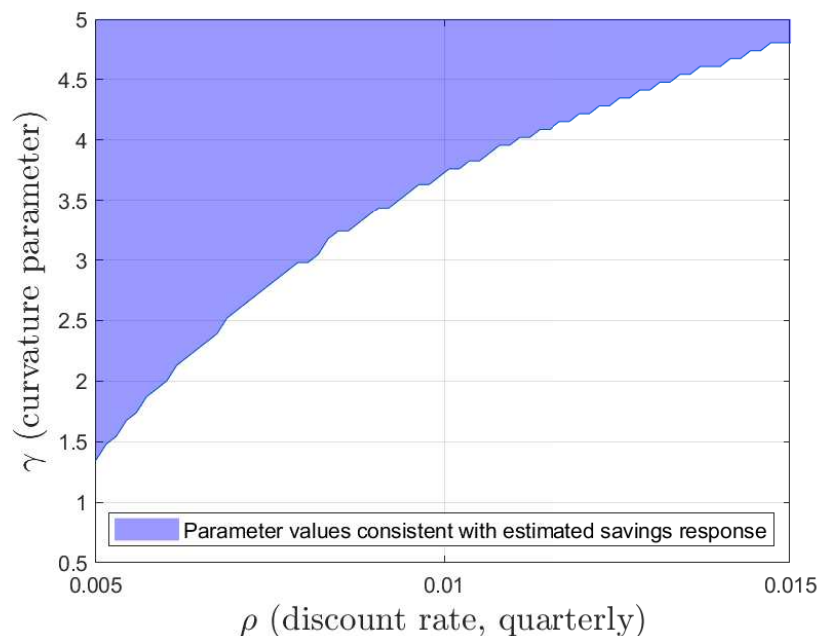


Figure 30: Set of data-consistent $\{\gamma, \rho\}$.

bound on γ increases as well. When households are discounting the future at a high rate, a large value of the curvature parameter is required to get a savings effect in line with our estimate. For relatively reasonable values of household's discount rate, the implied coefficient of relative risk aversion lies roughly between 1-4. This is consistent with macro-estimates of the coefficient of constant relative risk aversion, but inconsistent with micro-estimates which generally implies a curvature-parameter below 1, see for instance Holt and Laury (2002).

Appendix E: Model

The output of the tradable good $y_{m,T}$ and the non-tradable good $y_{m,N}$ in each municipality is produced using only labor as input

$$\begin{aligned} y_{m,T} &= \tilde{e}_{m,T} \\ y_{m,N} &= \tilde{e}_{m,N} \end{aligned}$$

where $\tilde{e}_{m,i}$ is the total employment in sector i and municipality m . Perfect labor mobility implies that wages and prices are equal across sectors.

Total labor supply to the tradable and non-tradable sector is given by

$$\tilde{e}_{m,T} + \tilde{e}_{m,N} = E_m \quad (28)$$

Total nominal demand from non-oil workers is assumed to be some amount given by $\mathcal{D}_{m,no}$. The non-oil workers have Cobb-Douglas preferences

$$U_{m,no} = C_{m,N,no}^\alpha C_{m,T,no}^{1-\alpha} \quad (29)$$

which yields demand functions

$$P_N C_{m,N,no} = \alpha \mathcal{D}_{m,no} \quad (30)$$

$$P_T C_{m,T,no} = (1 - \alpha) \mathcal{D}_{m,no} \quad (31)$$

Each municipality has a measure of f_m price-taking oil firms, which produce using intermediate inputs and labor

$$y_O = X_N^{a_N} X_T^{a_T} e_O^{a_L} \quad (32)$$

with $a_N + a_T + a_L = 1$. Market clearing in the oil sector is given by

$$y_O \leq \bar{y} \quad (33)$$

The combination of a Cobb-Douglas production function and the market clearing condition yields the following conditional factor demands

$$X_T = \frac{a_T}{P_T} \bar{y} \quad (34)$$

$$X_N = \frac{a_N}{P_N} \bar{y} \quad (35)$$

$$e_O = \frac{a_L}{w_O} \bar{y}$$

Oil workers have one unit of labor which they supply inelastically. Their income is therefore $\mathcal{D}_o = a_L \bar{y}$. We assume that oil workers have Cobb Douglas preferences

$$U_O = C_{m,N,o}^\alpha C_{T,o}^{1-\alpha} \quad (36)$$

which yields demand functions

$$P_N C_{N,o} = \alpha \mathcal{D}_o \quad (37)$$

$$P_T C_{T,o} = (1 - \alpha) \mathcal{D}_o \quad (38)$$

The following market clearing conditions pin down the prices P_N and P_T

$$y_{m,N} = C_{m,N,no} + \ell_m C_{N,o} + f_m X_{m,N} \quad \forall m \quad (39)$$

$$\int_{m \in \mathcal{M}} y_{m,T} dm = \int_{m \in \mathcal{M}} C_{m,T,no} dm + \bar{\ell} C_{T,o} + \bar{f} X_T \quad (40)$$

where $\bar{\ell} \equiv \int_{m \in \mathcal{M}} \ell_m dm$ and $\bar{f} \equiv \int_{m \in \mathcal{M}} f_m dm$, i.e the average fraction of oil-workers and non-oil workers across all municipalities.

In order to capture the possibility that oil workers crowd out employment in other sectors, we assume

$$\tilde{e}_{m,i} = e_{m,i} + \ell_m \kappa_i (1 - e_O) \quad (41)$$

Baseline equilibrium Armed with the demand functions in equations (30), (31), (34), (35), (37), and (38), the market clearing conditions (28), (39) and (40), we can solve for the equilibrium allocation.

First, inserting the production function and the factor demand equations into the market clearing condition for non-tradables yields

$$\tilde{e}_{m,N} = \frac{\alpha \mathcal{D}_{m,no}}{P_N} + \ell_m \frac{\alpha \mathcal{D}_o}{P_N} + f_m \frac{a_N}{P_N} \bar{y} \quad (42)$$

Similarly, for the tradadable good, we get

$$\int (E_m - \tilde{e}_{m,N}) dm = \frac{1 - \alpha}{P_T} \bar{\mathcal{D}}_{no} + \bar{\ell} \frac{1 - \alpha}{P_T} D_o + \bar{f} \frac{a_T}{P_T} \bar{y} \quad (43)$$

Combining the two yields

$$\int \left(E_m - \frac{\alpha \mathcal{D}_{m,no}}{P_N} - \ell_m \frac{\alpha \mathcal{D}_o}{P_N} - f_m \frac{a_N}{P_N} \right) dm = \frac{1 - \alpha}{P_T} \bar{\mathcal{D}}_{no} + \bar{\ell} \frac{1 - \alpha}{P_T} D_o + \bar{f} \frac{a_T}{P_T} \bar{y} \quad (44)$$

Integrating out the left-hand side, and using the fact that $P_{m,N} = P_T \forall m$, we get

$$\bar{E} - \frac{\alpha}{P_T} \bar{\mathcal{D}}_{no} - \bar{\ell} \frac{\alpha \mathcal{D}_o}{P_T} - \bar{f} \frac{a_N}{P_T} \bar{y} = \frac{1 - \alpha}{P_T} \bar{\mathcal{D}}_{no} + \bar{\ell} \frac{1 - \alpha}{P_T} D_o + \bar{f} \frac{a_T}{P_T} \bar{y}$$

which pins down the equilibrium price for non-tradable and tradable goods

$$P_T^* = P_N^* = \frac{\bar{\mathcal{D}}_{no} + \bar{\ell} D_o + \bar{f} (a_T + a_N) \bar{y}}{\bar{E}} \quad (45)$$

which yields equilibrium non-oil worker employment in the non-tradable and tradable sectors:

$$e_{m,N}^* = \frac{\alpha \mathcal{D}_{m,no}}{P_N^*} + \ell_m \frac{\alpha \mathcal{D}_o}{P_N^*} - \ell_m \kappa_N (1 - e_O) + f_m \frac{a_N \bar{y}}{P_N^*} \quad \forall m \quad (46)$$

and

$$e_{m,T}^* = \frac{1 - \alpha}{P_T^*} \bar{\mathcal{D}}_{no} + \bar{\ell} \frac{1 - \alpha}{P_T^*} D_o - \ell_m \kappa_T (1 - e_O) + \bar{f} \frac{a_T \bar{y}}{P_T^*} \quad \forall m \quad (47)$$

Oil shock Suppose now that prices are fixed at their baseline values and there is a shock to oil demand $d\bar{y}$. The change in the non-oil worker employment in the tradable sector is then given by

$$\frac{de_{m,T}}{d\bar{y}} = \frac{1 - \alpha}{P_T^*} \frac{\partial \bar{\mathcal{D}}_{no}}{\partial \bar{y}} + \bar{\ell} \frac{1 - \alpha}{P_T^*} \frac{\partial D_o}{\partial \bar{y}} + \ell_m \kappa_T \frac{\partial e_o}{\partial \bar{y}} + \bar{f} \frac{a_T}{P_T^*} \quad \forall m \quad (48)$$

Comparing two different municipalities m and m' , the cross-sectional difference in employment of non-oil workers in the tradable sector is given by

$$\frac{de_{m,T} - de_{m',T}}{d\bar{y}} = (\ell_m - \ell_{m'}) \kappa_T \frac{\partial e_o}{\partial \bar{y}} \quad (49)$$

The change in the employment of non-oil workers in the non-tradable sector is given by

$$\frac{de_{m,N}}{d\bar{y}} = \frac{\alpha}{P_N^*} \frac{\partial \mathcal{D}_{m,no}}{\partial \bar{y}} + \ell_m \frac{\alpha}{P_N^*} \frac{\partial \mathcal{D}_o}{\partial \bar{y}} + \ell_m \kappa_N \frac{\partial e_o}{\partial \bar{y}} + f_m \frac{a_N}{P_N^*} \quad (50)$$

Hence, cross-sectional difference in the employment of non-oil workers in the non-tradable sector is given by

$$\frac{de_{m,N} - de_{m',N}}{d\bar{y}} = \frac{\alpha}{P_N^*} \left(\frac{\partial \mathcal{D}_{m,no}}{\partial \bar{y}} - \frac{\partial \mathcal{D}_{m',no}}{\partial \bar{y}} \right) + (\ell_m - \ell_{m'}) \frac{\alpha}{P_N^*} \frac{\partial \mathcal{D}_o}{\partial \bar{y}} + (\ell_m - \ell_{m'}) \kappa_N \frac{\partial e_o}{\partial \bar{y}} + (f_m - f_{m'}) \frac{a_N}{P_N^*} \quad (51)$$

To simplify equation (51), assume that cross-municipal differences in demand from non-oil workers is driven by employment in the non-tradable sector. That is,

$$\frac{\partial \mathcal{D}_{m,no}}{\partial \bar{y}} d\bar{y} \approx w_N \frac{de_{m,N}}{d\bar{y}} \quad (52)$$

where w_N is the non-tradable sector wage. Inserting this in to equation (50) and solving for $\frac{de_{m,N}}{d\bar{y}}$ yields

$$\frac{de_{m,N}}{d\bar{y}} \approx \frac{1}{1-\alpha} \left(\ell_m \frac{\alpha}{P_N^*} \frac{\partial \mathcal{D}_o}{\partial \bar{y}} + \ell_m \kappa_N \frac{\partial e_o}{\partial \bar{y}} + f_m \frac{a_N}{P_N^*} \right) \quad (53)$$

Hence, the difference in non-tradable employment between two municipalities is given by

$$\frac{de_{m,N} - de_{m',N}}{d\bar{y}} = \frac{1}{1-\alpha} \left((\ell_m - \ell_{m'}) \frac{\alpha}{P_N^*} \frac{\partial \mathcal{D}_o}{\partial \bar{y}} + (\ell_m - \ell_{m'}) \kappa_N \frac{\partial e_o}{\partial \bar{y}} + (f_m - f_{m'}) \frac{a_N}{P_N^*} \right) \quad (54)$$