The Worker-Level Impact of Resource Sector Shocks*

Tazia Khushboo Jean-William Laliberte

November 24, 2025

(PRELIMINARY DRAFT. Click here for the most recent version.)

Abstract

We study the impact of the 2014-16 oil price crash on workers employed outside the oil and gas (O&G) sector in Canada using matched employer-employee data. We create three exposures through which the shock propagated to workers across industries and provinces: the labor market channel, the supply chain linkage, and the local demand channel. Workers employed at firms that previously lost many workers to O&G experienced a negative shock to their outside option due to displacement of their past co-worker networks. Firms that supply to oil and gas or depend on local demand from O&G workers experienced a drop in their sales due to the sector's collapse. Workers employed at such firms were exposed via decline in their share of quasi-rents. We find that the input-output and labor market channels are more important in explaining workers income losses: a standard deviation increase in exposure through each channel led to a 2% drop in incomes. Income losses through the local demand channel were negligible, about a 0.5% decline. Workers most affected belonged to the lower end of the skill distribution and were employed at the lowest-paying firms before the price shock.

* Khushboo: University of Calgary, tazia.khushboo@ucalgary.ca. Laliberte: University of Calgary and IZA, jean-william.lalibert@ucalgary.ca. We thank David Green and Trevor Tombe for helpful conversations. Funding from the Dunfield Fellowship is gratefully acknowledged.

1 Introduction

The oil and gas sector is the highest-paying sector in the Canadian economy. The oil price crash beginning in late 2014 was a significant blow to the revenues of the sector that triggered a recession in the extractive resource province of Alberta. Figure 1 illustrates how employment and earnings in the oil and gas sector are tied to the oil price.[1] The obvious outcome of this shock was massive unemployment in the oil and gas sector much larger than that experienced during the Great Recession. The oil and gas sector employment kept expanding during much of the Great Recession, with a massive decline experienced in 2009, when employment was about 80 percent of the 2013 level. Nonetheless, the drop in oil price beginning 2014 triggered a larger sectoral displacement lasting two years. Panel C shows that those displaced represented the lower spectrum of the skill distribution, as the average skill premium within the sector increased over the shock duration. Hence, the decline in the sector's employment earnings as shown in Panel B was largely driven by significant drops in the oil and gas firms' pay premiums to workers. The co-movement of employment and earnings gives a transparent picture of the extent to which oil price dictates workers' outcomes in the oil and gas sector. However, little is known about the extent to which this price shock propagated to workers in other sectors of the Canadian economy, which workers were most affected, and why.

[1] The West Texas Intermediate (WTI) crude oil futures contract traded on the New York Mercantile Exchange

Figure 1: Employment and Earnings in the Oil and Gas Sector Amid the Oil Price Shock.

Notes. Panel A shows employment in the oil and gas sector relative to the year prior to the oil price shock.

In this study, we examine how the 2014–16 oil price shock spills over to non-oil and gas workers in the economy through three different channels of exposure: the labor market channel, the supply chain channel, and the local demand channel. The shock affects workers based on their firm and province of employment, and hence, this study enables estimating the passthrough of firm-level shocks to worker outcomes using detailed employer-employee data. Our empirical strategy allows us to examine the relative strengths of the different channels in inducing worker impacts, and thus, helps identify which workers are most affected due to the price shock. While all three channels of exposure to the oil price shock led to income losses, those impacted through the supply chain channel exhibit are most affected. We find that a one standard deviation increase in shock exposure through the supply chain channel leads to a 2 percent decline in total income for non-oil and gas workers stably employed in the pre-shock period. We also document the direct impacts on oil and gas workers to put the indirect effects into perspective. Although the total income losses of direct workers can be 5 times as large as those of workers indirectly affected, the oil and gas sector employs a much smaller proportion of the Canadian workforce. Hence, the spillover effects can be large when aggregated at the total number of affected individuals. Note that, since we focus on tracking the outcomes of workers who had been stably employed prior to the shock, our estimates are a lower bound for the impacts on workers. Finally, we document the shock's impacts across the joint distribution of workers and firms. We find that the bulk of the impacts through the supply chain linkage is driven by the least skilled workers working in the least-paying firms.

Our results contribute to a large literature examining the labor market impacts of the forces of globalization and technological improvements. Much of the literature focuses on how China's meteoric rise as a dominant world exporter impacted the structure of employment and workers' outcomes in developed economies. Autor et al. (2013) spearhead this literature and estimate the worker displacement at the local labor market level. Acemoglu et al. (2016) augment their work by adding a macro focus and discussing two channels that could confound the estimates at the national level. However, our work is more closely related to studies by Autor et al. (2014) and Utar (2018). These worker-level studies improve substantially over the previous local labor market estimates by tracking the same set of workers longitudinally and controlling for endogenous worker movement across labor markets. In contrast, our study uses an exogenous

price crash in the global energy market and examines its impact on workers not only in the energy sector, but the spillover impacts reverberating across other sectors.

Our work improves upon previous studies in several ways. Studies on economic shocks such as the China Shock often focus on documenting the impacts on local labor markets (Autor et al., 2014; 2016). Our detailed employer-employee data enables us to examine shock impacts at the worker level. We track a set of workers attached to the same employer-province dyad in the pre-shock period and document what happens to their employment earnings, total income, and a host of other outcomes. Hence, our estimates are not affected by migration patterns that the shock might induce across local labor markets. Studies demonstrating indirect exposure based on the supply chain linkage typically use detailed firm-to-firm transactions (Dhyne et al., 2022;); however, most such analysis is restricted to examining firm-level impacts. In our case, we rely on publicly available provincial input-output tables to document the supply chain linkage enabling resource sector shock propagation. While this is a coarser measure than firm-level transaction data, this measure does provide us with sufficient variation across within industry-province cells. Moreover, our work juxtaposes the importance of two types of networks – supply chain linkages and employment networks – highlighting the relative strengths of each in inducing worker-level impacts.

This study is closely related to the literature on how impact of multinational corporations (MNCs) impact worker pay at domestic firms. MNCs, which are generally more productive than domestic firms, appear to have positive indirect effects on wages of workers located in the same local labor market (Timm, 2025; Setzler and Tintelnot, 2021) or working at enterprises that supply inputs to multinationals (Alfaro-Urena et al., 2021). In a similar way, the oil and gas sector is the most productive in the country. Nonetheless, a key difference is that while studies of MNE focus on idiosyncratic firm-level shocks, we focus on a much larger, sector-level shock which has the potential to severely disrupt an economy's industrial structure.

Our work also speaks to the sparse literature informing the adjustment process of workers vulnerable to the green energy transition. Much of this literature documents how the decline of the coal industry since the 1980s had adverse consequences for local employment and earnings (Black et al., 2005; Hanson, 2023). The coal industry does offer a unique case for learning about the labor market impacts of structural change. Nonetheless, it still represents an early phrase of the fossil fuel economy that has little relevance today. The O&G sector is the primary energy driver in today's world, and a substantial shift toward a green economy will necessarily involve pivoting away from this sector. Given this context, investigating even short-lived O&G sectoral

downturns can be of value for informing policy. In some sense, the episodic oil price shock we analyze provides a contemporary, realistic view into the short- and medium-term earnings losses and adjustment processes of an impending O&G phaseout. Furthermore, we document the impact on workers with indirect economic ties to the fossil fuel industry. These workers represent a much larger number of stakeholders than workers directly impacted by the green transition. Our results generate insights into which workers may be affected by a transition away from fossil fuels. For instance, Butikofer et al. (2025) show that the 1970s Oil boom in Norway fostered intergenerational mobility. Similarly, Haeck and Laliberte (2025) show that, in Canada, Oil and Gas workers are disproportionately drawn from low-income families, thereby contributing to intergenerational mobility. In this paper, we demonstrate that workers outside of the sector are also strongly affected by natural resource busts.

This paper is organized as follows. Section 2 describes the data sources and explains variable construction. Section 3 outlines the empirical strategy for evaluating the impacts of the oil price shock on workers. Section 4 provides descriptive evidence on the oil price shock's impacts on workers' earnings [TK]. Section 5 details the impacts on oil and gas workers. Section 6 sets out the main results for non-oil and gas workers. Section 7 concludes.

2 Data

Data Sources. Our main data source is comprised of confidential administrative tax information from the Canadian Employer-Employee Dynamics Database (CEEDD) over 2001-19. The oil price shock affects workers through its impact on their employers. As such, observing who works for which employer is a natural prerequisite for our analysis. The CEEDD's matched employer-employee longitudinal data enables us to track worker-firm relationships and enables us to observe detailed measures of firm and worker outcomes.

The T4 Record of Employment (T4ROE) in the CEEDD includes the universe of all T4 slips and records of employment issued by an employer for each of their workers every year, and hence, is the primary source for assigning workers exposure to the oil price shock. Note that a worker may hold multiple jobs during the same year. Hence, for each worker we pick the employer that issued the T4 with the highest earnings as the primary employer for that year. This allows us to create a unique worker-firm relationship for each year that represents the worker's actual labor market position compared to part-time or temporary jobs held. Nonetheless, the T4 earnings from this primary employer may not represent the worker's true

annual earnings if they switch jobs in a given year. Hence, we assign each worker the sum of T4 earnings across all jobs held during the year to capture total employment earnings.

The T1 Family Files (T1FFs) in the CEEDD contain information for the universe of all tax filers in Canada. From this data source, we obtain specific components of an individual's total income in addition to demographics such as age, gender, and place of residence on December 31. Since employers are legally required to issue T4s to all workers they employ, we observe T4 earnings regardless of individuals' tax filing behavior. We impute missing (pre-tax) total, and employment incomes for workers in a given year from their (non-missing) employer-issued T4 earnings. We assign a zero value to any further missing cases for income variables. All income variables are deflated to constant 2012 dollars using the national CPI. These administrative data do not include hours of work at the individual level, and hence, we our analysis will focus on income rather than wage effects of shock exposure.

The National Accounts Longitudinal Microdata File (NALMF) contains the 4-digit NAICS industry that each firm belongs to. Using this variable, we define the O&G sector to include firms in O&G extraction (NAICS4 2111) and support activities for mining, and O&G extraction (NAICS4 2131). In addition, the NALMF provides information on firm outcomes such as total revenue that we use to demonstrate firm-level impacts of the oil price shock as a first stage for the worker impacts.[2]

We complement this rich data with publicly available provincial input-output-tables to create our measure of supply linkage exposure to the oil price shock at the industry-province level. While our measure of supply linkage is coarser than that of studies using firm-to-firm transactions, it does constitute a substantial improvement compared with an industry-level measure based on national IO tables. Combining spatial and industry variation in supply chain relationships allows us to isolate the role of input-output linkages from that of national industry-specific shocks and of differential trends in economic outcomes across provinces.

Sample Selection. To plausibly attribute changes in workers' outcomes to the oil price shock, we focus on workers who had been stably employed during the pre-shock period. For all worker-

[2] Ideally, we would have preferred to examine firms' sales of good and services. However, our version of the CEEDD does not include this variable. We instead use the total revenue measure that accounts for gross sales as well as other forms of income such as grants, subsidies, and incentives from the government. Firms more affected by the oil price shock are generally more likely to receive government subsidies and grants. However, we do not expect firms to have unusually high levels of incentive receipt in a given year relative to their normal levels or the industry average in that year. Hence, we use firm and industry-year fixed effects to kill off the variation in total revenue coming from measurement error.

level analyses, we define an employer as a firm-province dyad to measure worker transition across firms and provinces with greater precision. The sample of stably employed workers include those who were no younger than 25 in 2010 and no older than 54 in 2013; had the same primary employer; and made at least $5,000 in (CPI-adjusted) annual income. We exclude persons whose year of death was in 2013.[3] This sample consists of approximately 5 million individuals. However, due to computational issues, we are unable to run the analyses on this full sample. Instead, we use a 20-percent stratified random sample to ensure the sector-specific employment composition carries over. We track the outcomes of this random sample of stably employed individuals over 2009-19. This means that we only include post-recession years in the pre-shock period. This ensures a plausible baseline for studying workers' adjustment following the oil price shock.[4] For our main analysis, we restrict the 20-percent sample to only non-O&G workers. Additionally, we use the sample of stable workers across all industries to estimate the direct effects of the price shock on O&G workers. We also use the full sample of workers aged 25-54 over 2001-2019 to estimates time-invariant AKM worker and firm fixed effects to explain mechanisms underlying workers' adjustment to the shock. We assign each stably employed worker in our sample a firm fixed effect based on their current primary employer in a given year.

Exposure Measure Construction. In our setting, firms in the highest-paying sector face a significant shock to their revenues. We study how this shock propagates to other sectors in the economy. To do so, create three measures of exposure to the oil price shock: the labor market exposure, the supply chain exposure, and the local demand exposure. All three measures exhibit spatial and temporal variation.

The labor market exposure or bargaining channel ($LM_{j(i)}$) is a measure of how vulnerable an employer is to losing their workers to the O&G sector. Past co-workers who are currently employed in the O&G sector can provide the firms' existing workers with information on job

[3] Individuals who have been stably employed over 2010-13 and whose year of death is 2013 might have smaller than regular earnings in that year based on the timing of death.

[4] Younger workers who are yet to formally enter the labor force would appear to have artificially low earnings in years further away from the sample window of 2010-13. Similarly, earnings of workers during the recession years would be abnormally lower than under usual labor market conditions. Hence, we only include post-recession years in our analysis sample to avoid depressing the pre-shock earnings and to prevent attenuation of the difference-in-difference estimates toward zero. Note that our post-shock event study estimates would not affected by the inclusion of these years since we use 2013 as the baseline year. However, the event study estimates for years further away from 2010 would illustrate upward sloping pre-trends that are not true anticipation effects.

prospects and act as job references (Caldwell and Harmon, 2022). Hence, workers employed at firms that have lost many of its workers to the O&G sector can more credibly threaten to leave as well (Green et al, 2023). This gives the workers enhanced bargaining leverage with their current employers as the value of their outside options tends to reflect the value of employment in the highest-paying O&G sector. We create the measure as follows: $LM_j = \frac{\sum_{2006}^{2012} N_{j,O\&G,t+1}}{\sum_{2006}^{2012} N_{j,t}}$. Each firm j in year t employs $N_{j,t}$ number of workers. $N_{j,O\&G,t+1}$ tracks how many workers move to firms in the O&G sector in year $t+1$. Hence, the proportion $\frac{N_{j,O\&G,t+1}}{N_{j,t}}$ is the probability that firm j's workers will leave for the O&G sector in year $t+1$. We make this probability time-invariant by summing the total number of moves across all years $t+1$ and express this quantity as a proportion of the sum of initial workers employed across all years. This gives us the probability that a worker initially employed at firm j transitions to the O&G sector in any given year.

We construct the supply linkage exposure using provincial input-output (IO) tables. For each industry k in province p, we calculate the share of sales that go directly and indirectly to the O&G sector (Long and Plosser, 1983; Acemoglu et al., 2012; Acemoglu et al., 2016). Indirect shares represent the share of sales to other industries which then sell to the O&G sector. The supply linkage exposure ($IO_{j(i)}$) is the sum of these direct and indirect industry-province level sales to the O&G sector. We assign workers this exposure based on their pre-shock industry and province of employment.

To measure shock exposure through local demand channels, in the spirit of Mian and Sufi (2014), we combine cross-sectional spatial variation in the concentration of O&G workers with variation in tradability across industries. Firms in locations where O&G workers are the primary spenders in the local economy would be harder hit by the oil price shock. The CEEDD does not record granular location information on firms. However, it does provide synthetic data on individuals' Census Division (CD) of residence, which offer far greater granularity than the province of the firm's location. Hence, we exploit this finer spatial variation at the employee-level to construct our local demand exposure. Among all individuals living within a CD during the pre-shock period, we compute the share of employment earnings of O&G workers that gives us a magnitude of the local spending accruing to O&G workers. In addition, at the province level, we compute the share of sales for each industry that are out-of-province (including inter-province sales and international sales) to measure the tradability of each province-specific industry. One minus this tradability measure gives the share of province-specific industry output that is consumed locally. We obtain the local demand exposure (LD_i) as a product of the share of

O&G earnings times the share of the province-industry output that is locally consumed. We then average this measure across employees to obtain a time-invariant, firm-specific local demand exposure.

3 Identification Strategy

The primary focus of this study is to document the spillover impacts of the oil price shock on workers in the non-oil and gas sector. We compare the outcomes of more-versus-less affected non-oil and gas workers based on how exposed they are to the oil price shock through each of the channels using the following specification:

$$y_{it} = \alpha_i + \gamma_t IO_{j(i)} + \delta_t LM_{j(i)} + \theta_t LD_{j(i)} + \lambda_{g(i),t} + \epsilon_{it}$$

Here, y_{it} is worker i's outcome for year t who worked at firm j during the pre-shock period of 2010-13. We control for worker fixed effects as well as baseline industry-by-year and baseline employment province-by-year fixed effects.

To put the spillover effects into perspective, we also document the impacts of the price shock on oil and gas workers. The oil and gas workers are the highest paid workers in the country, and this precludes finding a natural counterfactual group of workers employed in some other industry. One challenge that we face to estimating the direct impacts of the shock is to define a credible counterfactual group of workers. Previous studies have compared the outcomes of observationally similar workers who differ only in terms of their exposure to a negative shock or policy using selection-on-observables (e.g. LMW, 2020; Kovak and Morrow, 2024). We attempt to improve upon this approach by using synthetic difference-in-difference (DID) (Arkhangelsky et al., 2021). This is a data-driven method that uses time and industry weights to create the counterfactual path of the outcome for O&G workers instead of controlling for covariates based on researcher discretion or previous studies. The method uses industry and year weights in the sum of squared error minimization to estimate the causal impact of the oil price shock on earnings

4 Direct Effects on Oil and Gas Workers

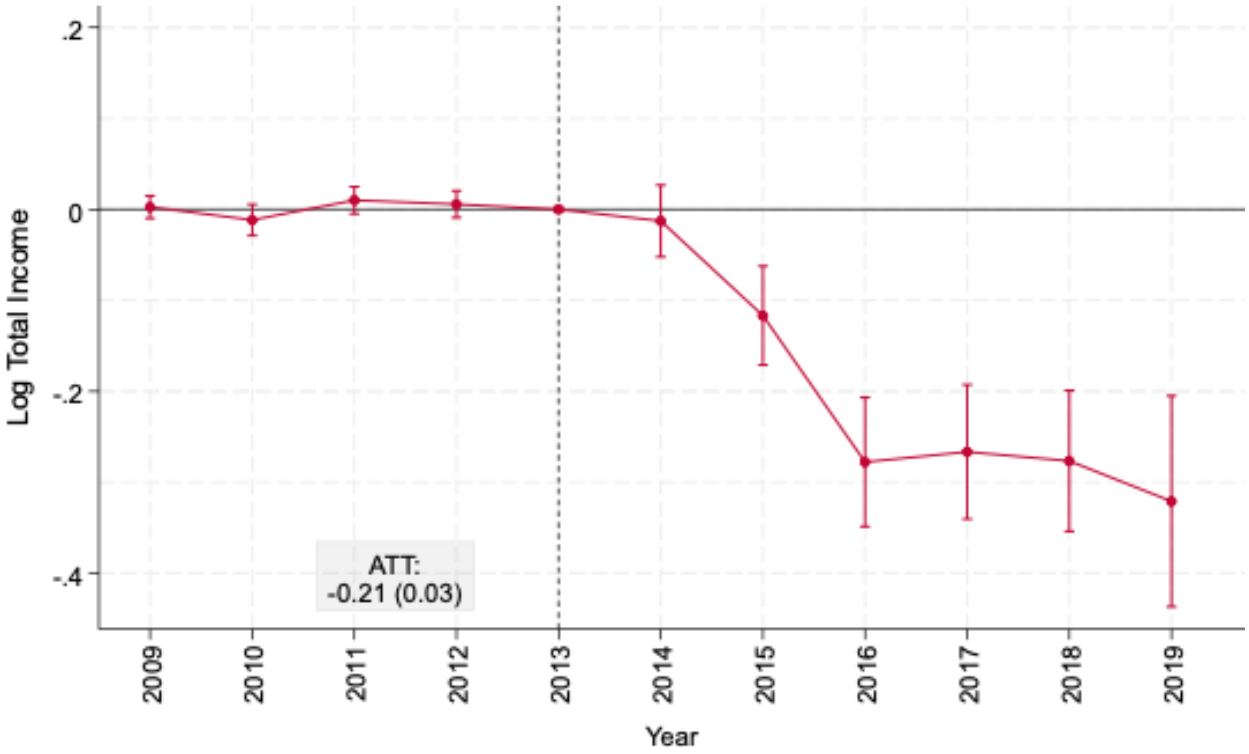

Figure 2: Direct Impact of the Oil Price Shock on Workers

Notes. Capped bars indicate 95-percent confidence intervals. Total income is workers' total income as reported in the T1FF in the CEEDD database.

We begin with estimating the impact of the oil price shock on O&G firms and workers using synthetic DID. Figure A1 illustrates the impact of the shock on firms. Surprisingly, the oil price shock does not seem to have a significant adverse impact on O&G firms' total revenue. It is important to note that our measure of total revenue includes other revenues such as investment and government incentives along with sales. As such, any negative impacts on sales could be counteracted by positive revenues from other sources. Nonetheless, Figure 2 as well as our descriptive plots in Figure 1 do provide clear, substantive evidence of adverse consequences of the shock for O&G workers. We find that the synthetic DID weights created a counterfactual income path for O&G workers that was similar to their actual income path prior to the price decrease. However, total income starts to drop as the oil price takes a hit, with the decline becoming progressively larger. By 2016, O&G workers experience a massive income decline of 24 percent. Most importantly, these income losses do not show any signs of recovery over the next three years. Together, these results suggest that firms passed on the oil price shock to their workers, leading to substantial income losses for O&G workers. Crucially, this result points

to the career disruption faced by the employment networks of workers exposed to the oil price shock through the labor market channel.

5 Worker-level Spillover Effects of the Oil Price Crash

We move on to the indirect effects from exposure to the oil and gas sector. Workers exposed to the shock through all three channels experience some kind of income losses. A one standard deviation increase in shock exposure leads to a 2 percent income loss at the peak of the price shock. Overall, workers exposed through the supply linkage and the labor market channel are the most affected, with the impact on workers affected through the local demand channel being about half as much. This illustrates that employment networks are as just as important as production networks for sectoral shock propagation when the collapsed sector is a high-paying employer.

We also document the importance of exposure intensity to compare with our direct effects, using a static ventile analysis as shown in Figure 2A. Workers across the distribution of exposure through the supply linkage channel is experience negative effects of similar magnitudes, indicating a global effect. However, exposure through labor market linkages kicks in only after crossing a certain threshold. Nonetheless, the distribution of effects through the local demand channel is rather nonlinear, with income losses becoming significant only in the last decile. Overall, while exposure intensity matters, marginal income losses are mostly stable across the exposure distributions, we stick to the linear specification for clarity of exposition. Overall, we find that the most affected workers through the input-output linkage channel experience income losses of about 5 percent, those affected through the bargaining channel experience losses of about 4 percent, and those affected through the local demand channel experience losses of about 2 percent, compared to the 24 percent persistent losses faced by workers in the O&G sector.

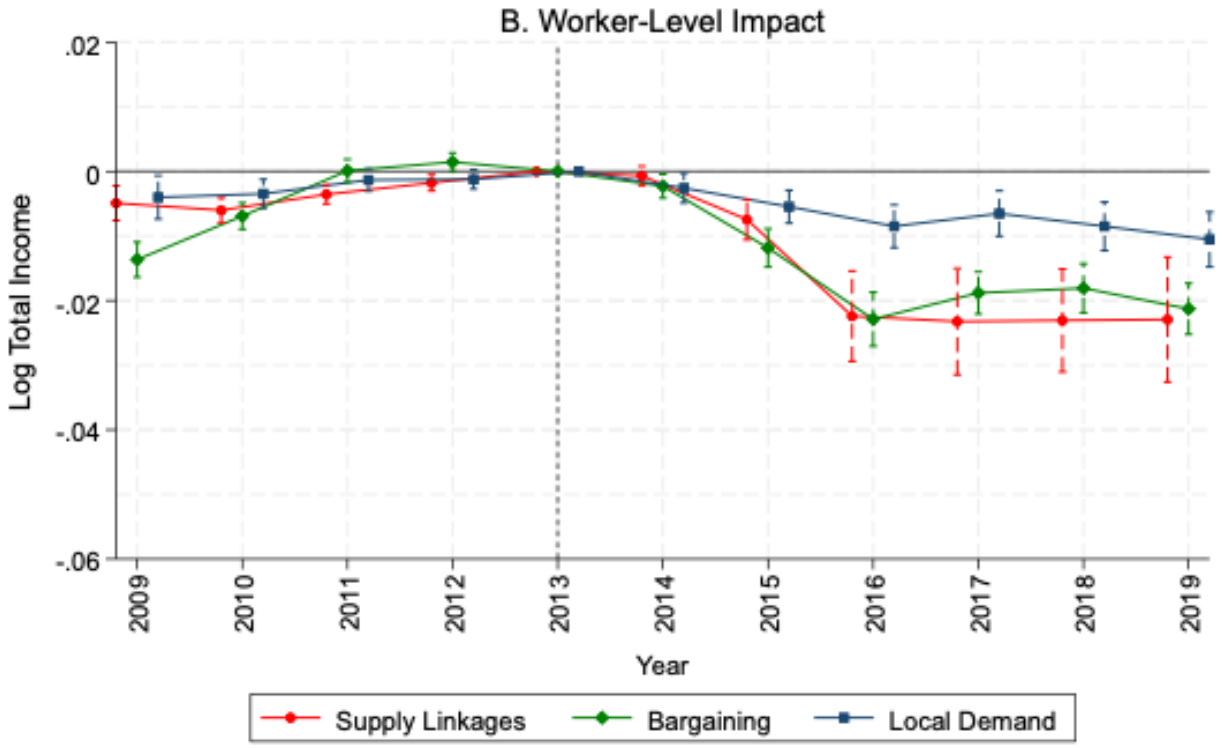

Figure 3: Impacts of Exposure to the Oil Price Shock Through Distinct Channels.

Notes. Capped bars indicate 95-percent confidence intervals. Total income is workers' total income as reported in the T1FF in the CEEDD database.

Our generalized econometric model gives us the advantage of comparing the relative strength of exposure through all three channels at the same time. However, previous studies mainly focus on a single channel independently. Hence, we also run separate regressions of the log of total income on each exposure measure using the same fixed effects and standard error clustering to allow comparisons with the estimates from our generalized model. Figure 4 shows that the point estimates across all three exposure measure is slightly larger in magnitude than that in our generalized specification. Nonetheless, the overlapping confidence intervals suggest that the differences are not statistically significant. As such, our results signal that the bias from considering a single exposure channel is limited.

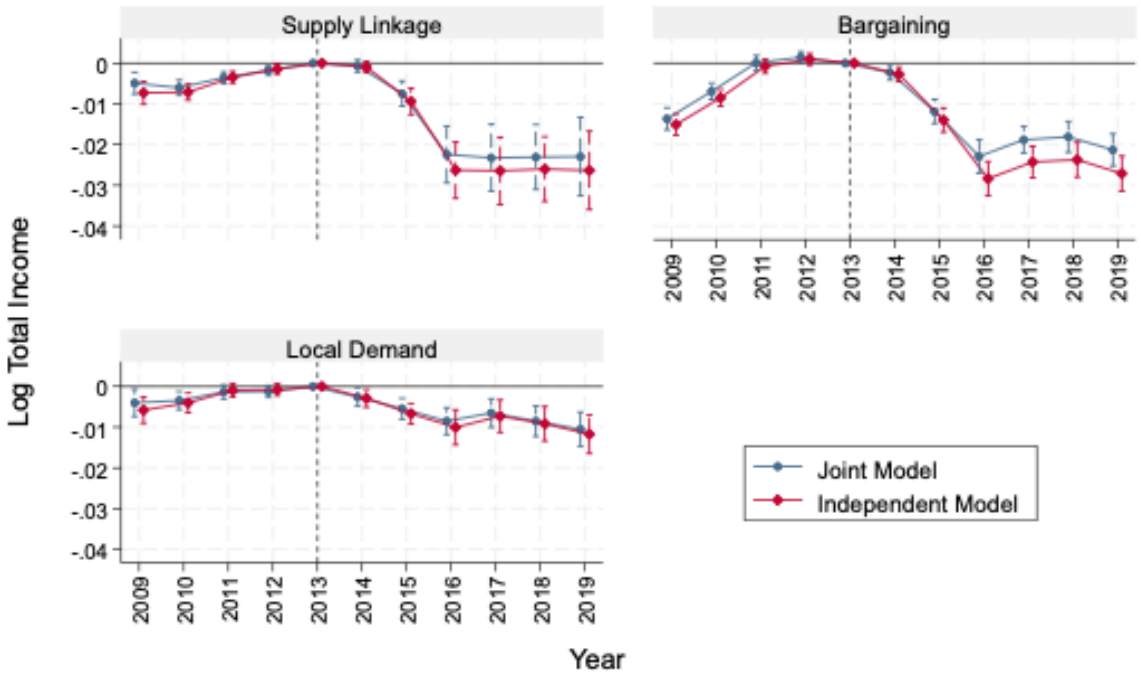

Figure 4: Robustness Checks: Considering Each Exposure Channel on Its Own

Notes. Capped bars indicate 95-percent confidence intervals. Total income is workers' total income as reported in the T1FF in the CEEDD database.

Figure 5 reruns the main specification excluding workers initially employed in Alberta firms. We see that the effects are stronger for workers initially employed in Alberta who are affected through input-output linkages and employment networks. However, workers outside of Alberta face stronger local labor demand effects. This speaks to the work of Green et al. (2019), Marchand (2015), and Fortin and Lemieux (2015) with regard to a negative shock on the O&G sector.

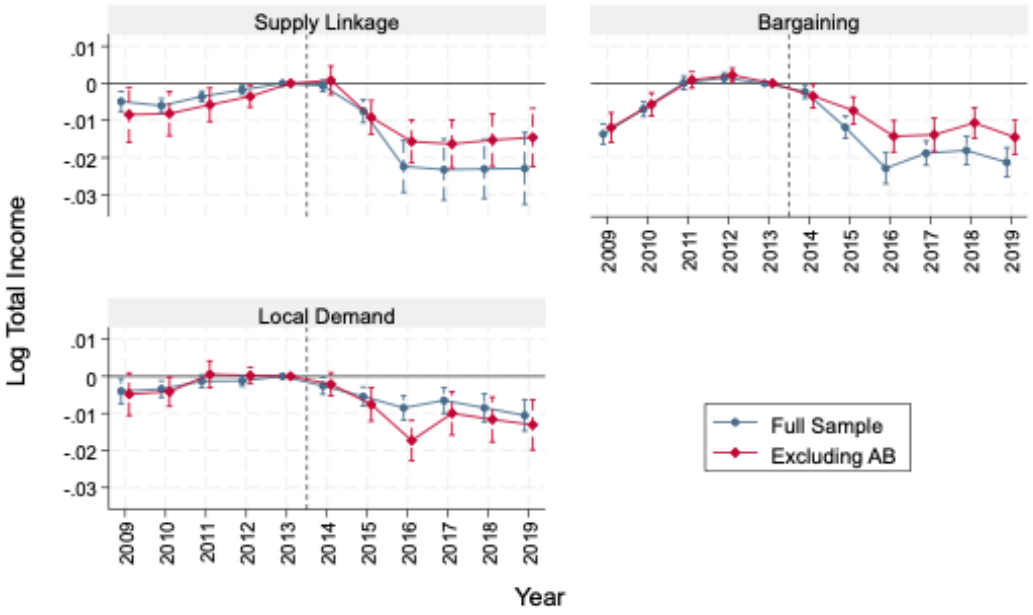

Figure 5: Robustness Checks: Excluding Alberta from the Analysis

Notes. Capped bars indicate 95-percent confidence intervals. Total income is workers' total income as reported in the T1FF in the CEEDD database.

In Figure 6 below, we decompose workers' dollar incomes into employment income, self-employment income, and a residual income and run the main regression specification. We find that, overall, non-O&G workers experience a drop in incomes of about $2,000, with over 90 percent of that loss originating from drops in employment income. Since we consider a sample of stably employed workers prior to the shock hitting the O&G sector, these income losses represent the lower bound. Workers who are in less stable employment relationships are more likely to experience higher income losses and disruptions due to exposure through the different channels. Note that workers experiencing lower self-employment incomes are mainly being affected through the local demand channel, although the effect is hardly tangible.

Figure 6: Decomposition Along Income Sources

Notes. Capped bars indicate 95-percent confidence intervals. Total income is workers' total income as reported in the T1FF in the CEEDD database.

Next, we consider heterogeneous impacts of exposure to the oil price shock along several dimensions. First, we show in Figure 7 that women bore slightly higher income losses than men, although the difference is not statistically significant. Nonetheless, in appendix Figure 8, we find that it is largely young workers who were aged at most 30 in the year prior to the price shock who drive the income losses through each of the exposure channels.

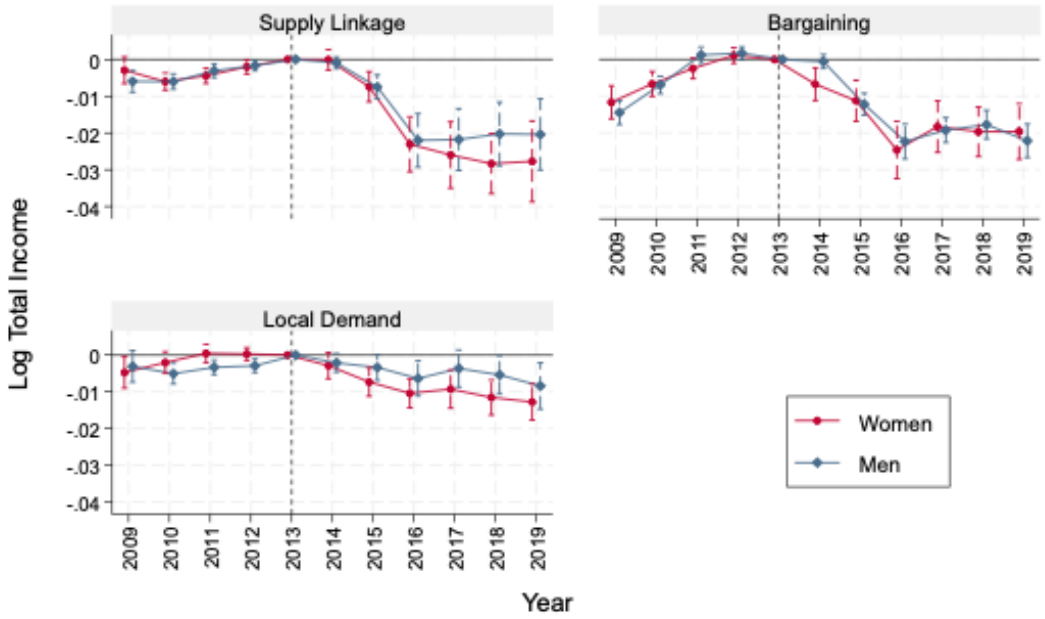

Figure 7: Heterogeneity Analysis by Gender

Notes. Capped bars indicate 95-percent confidence intervals. Total income is workers' total income as reported in the T1FF in the CEEDD database.

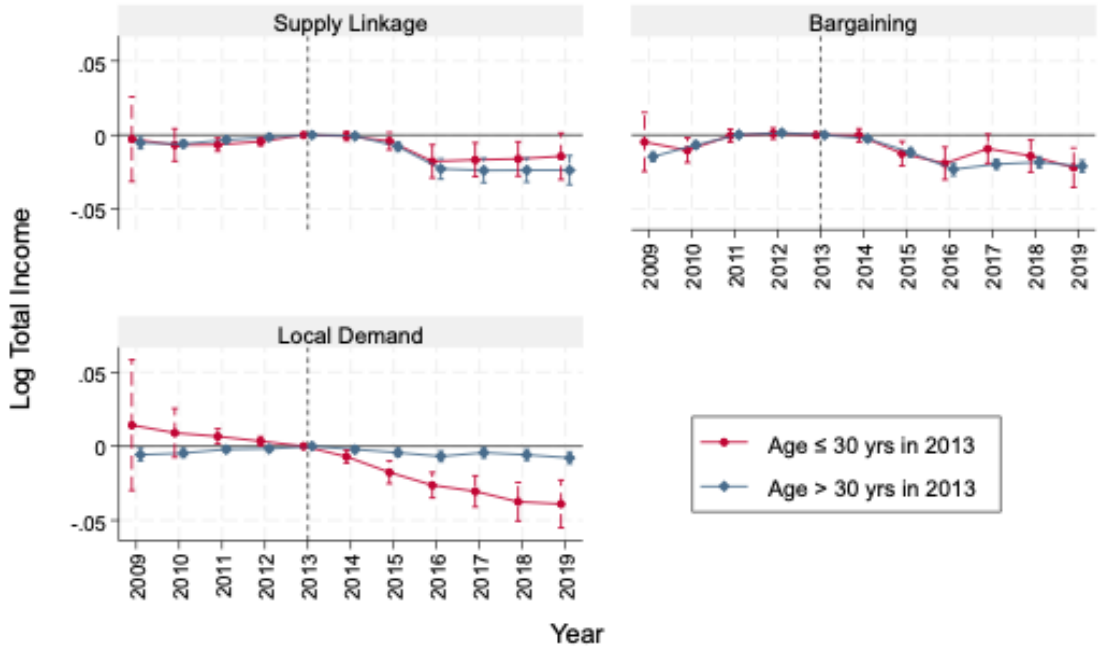

Figure 8: Heterogeneity Analysis by Age

Notes. Capped bars indicate 95-percent confidence intervals. Total income is workers' total income as reported in the T1FF in the CEEDD database.

Furthermore, we document heterogeneity by workers' pre-shock position in the skill premium and their firms' pay premium distributions. To do this, we conduct an AKM decomposition of workers' employment earnings into person and firm fixed effects using the pre-shock sample spanning 2001-13. Next, we summarize these distributions of fixed effects into quintiles. We assign each worker in our stable sample these fixed effect quintiles based on whether they showed up in the largest connected sample used in the AKM regression. Using our generalized specification, ee decompose the impacts of exposure through each channel by worker and firm fixed effects. The results of this granular heterogeneity analysis are shown in Figure 9 below. We find that the income losses through the supply linkage and bargaining channels are concentrated among workers in the lower part of the skill distribution who are employed by the low-paying firms. Interestingly, workers who are affected through the local demand channel are mostly medium-skilled workers initially employed at lower-paying firms. These results further help us understand the severity of exposure along the joint worker-firm distribution.

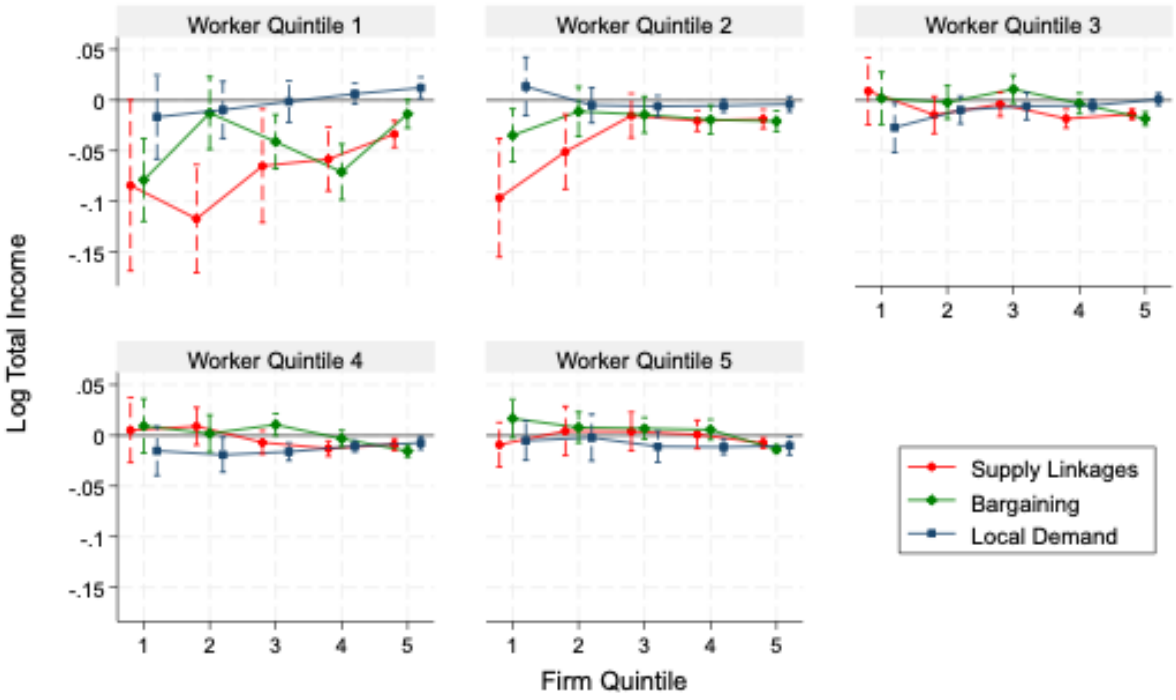

Figure 9: Heterogeneity Analysis by Age

Notes. Capped bars indicate 95-percent confidence intervals. Total income is workers' total income as reported in the T1FF in the CEEDD database.

6 Conclusion

In this study, we explore the various channels through a shock originating in a high-paying sector can ripple through the economy. Our setting is the oil price shock over 2014-16 in Canada. We consider three distinct channels through which workers can be connected to the oil price shock: the supply linkage channel, the labor market channel, and the local demand channel. In comparison to the massive income losses of 24 percent for directly affected workers, the most exposed workers outside the sector of immediate shock experience 5% to 2% losses. Nonetheless, the aggregate losses across workers in the economy is much larger since the oil and gas sector accounts for a very small proportion of workers in the Canadian economy.

7 References

Acemoglu, Daron, David Autor, David Dorn, Gordon H Hanson, and Brendan Price, "Import competition and the great US employment sag of the 2000s," Journal of Labor Economics, 2016, 34 (S1), S141–S198.

Alfaro-Urena, Alonso, Isabela Manelici, and Jose P Vasquez, "The effects of multinationals on workers: evidence from Costa Rican microdata," Technical Report 2021.

Autor, David H, David Dorn, Gordon H Hanson, and Jae Song, "Trade adjustment: Worker-level evidence," The Quarterly Journal of Economics, 2014, 129 (4), 1799–1860.

Barrot, Jean-Noel and Julien Sauvagnat, "Input specificity and the propagation of idiosyncratic shocks in production networks," The Quarterly Journal of Economics, 2016, 131 (3), 1543–1592.

Butikofer, Aline, Antonio Dalla-Zuanna, and Kjell G Salvanes, "Breaking the links: Natural resource booms and intergenerational mobility," Review of Economics and Statistics, 2025, 107 (2), 306–323.

Caldwell, Sydnee and Nikolaj Harmon, "Outside options, bargaining, and wages: Evidence from coworker networks," Unpublished manuscript, Univ. Copenhagen, 2019, pp. 203–207.

Card, David, Jesse Rothstein, and Moises Yi, "Industry wage differentials: A firm-based approach," Journal of Labor Economics, 2024, 42 (S1), S11–S59.

Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi, "Supply chain disruptions: Evidence from the great east japan earthquake," The Quarterly Journal of Economics, 2021, 136 (2), 1255–1321.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Toshiaki Komatsu, Magne Mogstad, and Felix Tintelnot, "Foreign demand shocks to production networks: Firm responses and worker impacts," Technical Report, National Bureau of Economic Research 2022.

Gathmann, Christina, Ines Helm, and Uta Sch ̈onberg, "Spillover effects of mass layoffs," Journal of the European Economic Association, 2020, 18 (1), 427–468.

Haeck, Catherine and Jean-William Laliberté, "Careers and Intergenerational Income Mobility," American Economic Journal: Applied Economics, 2025, 17 (1), 431–458.

Helm, Ines, "National industry trade shocks, local labour markets, and agglomeration spillovers," The Review of Economic Studies, 2020, 87 (3), 1399–1431.

Kovak, Brian K and Peter M Morrow, "The long-run labor market effects of the canada-us free trade agreement," Technical Report, National Bureau of Economic Research 2022.

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

Marchand, Joseph, "The distributional impacts of an energy boom in Western Canada," Canadian Journal of Economics, 2015, 48 (2), 714–735.

Mian, Atif and Amir Sufi, "What explains the 2007–2009 drop in employment?" Econometrica, 2014, 82 (6), 2197–2223.

Setzler, Bradley and Felix Tintelnot, "The effects of foreign multinationals on workers and firms in the United States," The Quarterly Journal of Economics, 2021, 136 (3), 1943–1991.

Statistics Canada, "Provincial Symmetric Input-Output Tables, 2014," 2018.

Timm, Lisa Marie, "The impact of multinational firms on local labour markets – what drives the effect on domestic workers?" 2025.

Appendix Figures

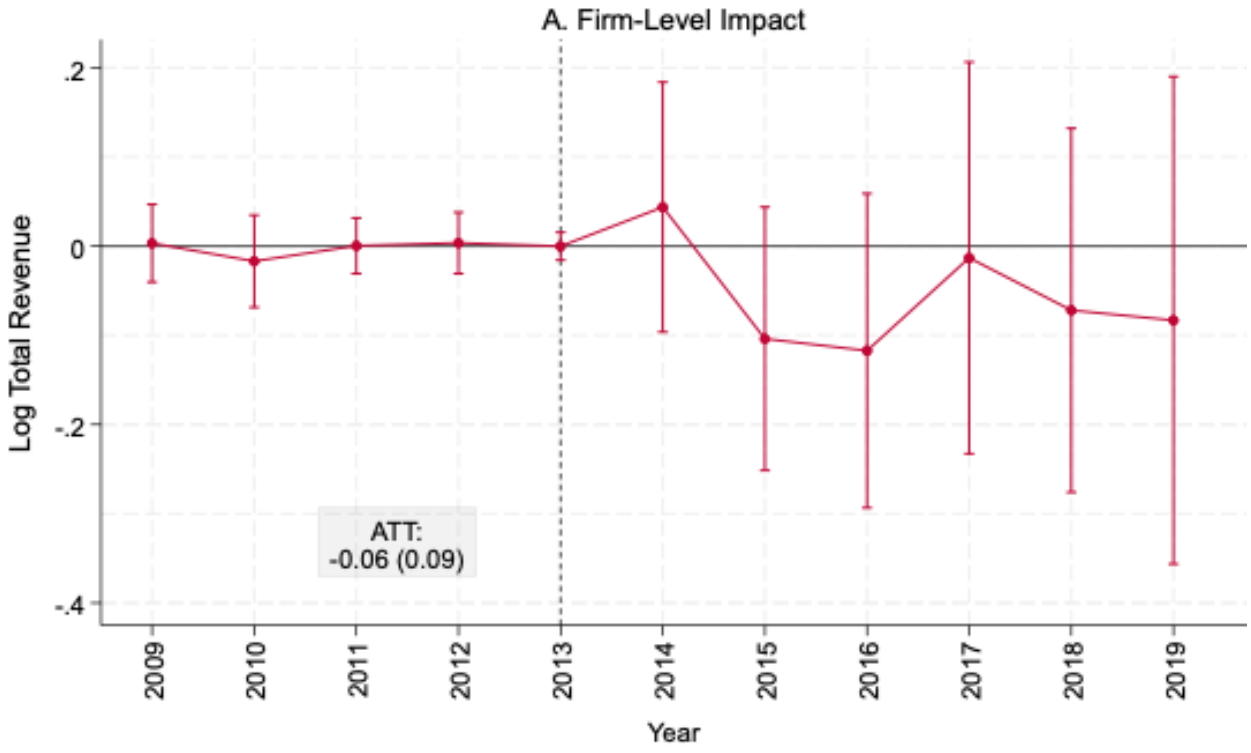

Figure A1: Direct Impact of the Oil Price Shock on Firms.

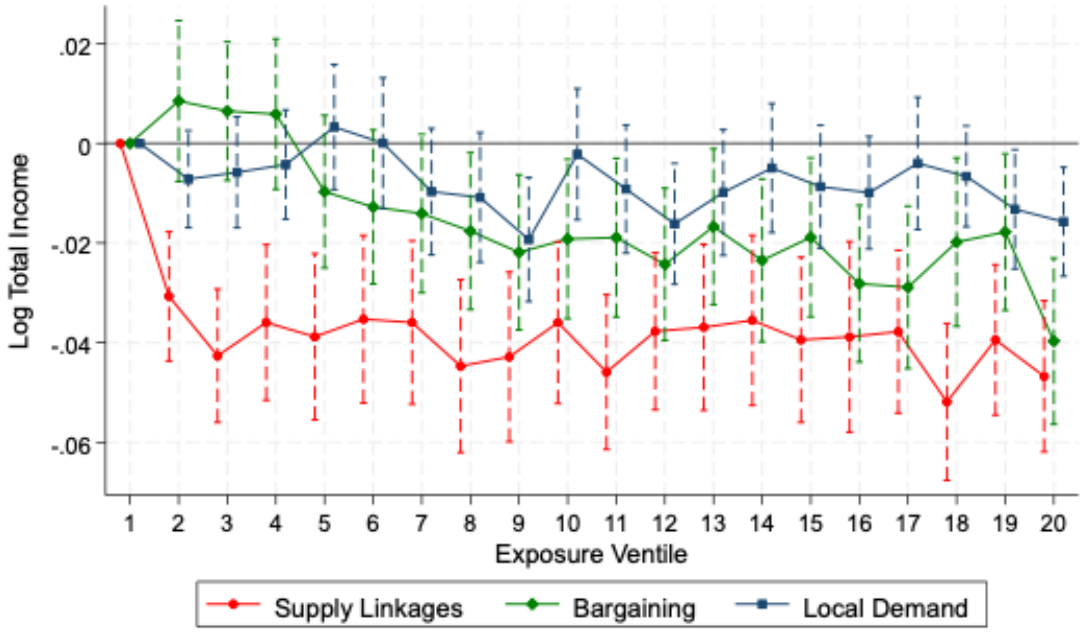

Figure A2: Impacts by Exposure Ventiles.