

Is It All Relative? The Health Impacts of Changes to Absolute and Relative Income *

Maripier Isabelle[†]

Boriana Miloucheva[‡]

May 2021

Abstract

This paper aims to disentangle absolute and relative income effects on health, with a focus on mental health. To identify the effect of a change of one's position in the income distribution, we exploit the heterogeneous effects of exogenous movements in the price of oil on the distribution of income. Using hospitalization records linked to census data, we find that oil workers who have many neighbors in the oil industry are less likely, and non-oil workers more likely, to seek hospitalization after oil prices rise. We examine a number of possible explanations for this finding. We show that it is unlikely to be due to a decline in these individuals' real purchasing power. It is consistent, however, with a relative income effect on health. Our results shed new light on mechanisms through which income inequality might affect people's well-being.

Keywords: Relative income, inequality, income shocks, health outcomes, mental health

JEL codes: I10, I14

*We would like to thank Arthur Silve, Michael Baker, Dwayne Benjamin, Shari Eli, Dylan Gordon, Mark Stabile, Kaisa Snellman, Mark Kelly, Jill Furzer and participants at the INSEAD, UQAM and Université Laval seminar series, as well as the 2019 Canadian Economics Association Meetings and International Health Economics Conference for invaluable feedback. Much of the analysis for this paper was conducted at the Toronto Region Statistics Canada Research Data Centre, which is part of the Canadian Research Data Centre Network (CRDCN). The services and activities provided by the CRDCN are made possible by the financial or in-kind support of the SSHRC, the CIHR, the CFI, Statistics Canada and participating universities whose support is gratefully acknowledged. Maripier Isabelle gratefully acknowledges support for this study from the Sentinel North Partnership Research Chair in Economics and Brain Health. Boriana Miloucheva acknowledges support for this study from the CRDCN Emerging Scholars Grant. The views expressed in this paper do not necessarily represent the views of the CRDCN or its partners. Comments can be sent to: maripier.isabelle@ecn.ulaval.ca

[†]Université Laval, Pavillon J.-A.-DeSeve, 1025 avenue des Sciences Humaines, Quebec (Canada) G1V0A6.

[‡]University of Toronto, 150 St. George St., Toronto, Ontario M5S 3G7, Canada

1 Introduction

Combined with a rapid concentration of income among highest earners, recent reversals in morbidity and mortality trends among certain groups in the U.S. have renewed interest in the relationship between health and the evolution of income distributions. The potential consequences of reduced social mobility on well-being among individuals who used to identify as working middle class is also fuelling policy conversations. Speaking to these questions, work spearheaded by Case and Deaton (2015) and Case and Deaton (2017) have recently suggested that the impression of *slipping down the social ladder*, or *falling behind*, might have been an important contributing factor to the stark increase in deaths attributable to suicides, prescription-drug overdoses and alcohol-related liver diseases since the late 1990s among middle-aged white non-Hispanic Americans without a college degree. While work on these *deaths of despair* primarily draws on American context and data, comparable phenomena have also been documented beyond the U.S (e.g. Baker et al. (2019), Baker et al. (2021)).

More generally, the relationship between relative income, health and subjective well-being indices has attracted significant attention in recent years.¹ However, empirically assessing the causal impact of changes to relative income or socioeconomic situation is challenging. First, both absolute and relative income are likely endogenous inputs in the health production function. For example, a person with chronic pain may be forced to limit her participation in the labour market, causing a reduction in earnings and a potential slip down the income distribution.² Second, disentangling the health impacts of changes in one’s relative income level from those of a change in her absolute level of income is not straightforward. Indeed, in general, changes in individuals’ income levels simultaneously affect their rank or position within the income distribution.

In this paper, we propose an empirical setting and approach to address these challenges in order to investigate how changes in an individuals’ absolute and relative income within a local distribution might affect their consumption of inpatient care services. To address the potential issues associated with reverse causality between health and income, we exploit movements in the price of oil and their differential impact on individuals’ earnings based on their industry of work. Indeed, changes in the price of oil disproportionately affect (with a lag) the earnings and employment of individuals working in the oil and gas extraction industry in the short-run, without having a similar impact on those working in other industries.³

These same movements in the price of oil also provide an opportunity to explore the specific health impacts of one’s movement along a reference income distribution. Indeed, the impact that a change in oil prices can have on local income distributions will vary in importance based on the share of individuals working in the oil industry in the corresponding local labour market. To separately identify the health impacts of changes to absolute and relative income, our empirical

¹See for example Clark and Oswald (1996) Clark et al. (2010), Daly and Wilson (2009), Daly et al. (2013), Ifcher et al. (2018), Brodeur and Fleche (2019).

²See Lleras-Muney (2018) for an interesting discussion of recent work looking at reverse causality in the relationship between income and health.

³A discussion on the relevance and exogeneity of these shocks is provided in section 3.

approach therefore exploits the various combinations of individuals' own labour market activity and of the industrial composition of their neighbourhood.

To implement this approach, we use 4 years of administrative hospital discharge records linked to the 2006 Canadian long-form census, a representative 20% sample of the country's population. In addition to providing longitudinal individual-level information on the nature and severity of diagnoses for all inpatient stays over that period, the linked data allows us to identify individuals' industry of employment, information that is typically not included in health records. The census data also contains information that enables us to identify each individual's neighbourhood of residence, which we use as a reference group. With this geographic information, we can also measure the concentration of oil workers within each neighbourhood, and the proportion of individuals within each reference group whose income should be impacted by changes in the price of oil.

While our research question is universal, and our results speak to various contexts, using Canadian data is critical to our empirical strategy. First, the importance of the oil and gas extraction industry across local labour markets in the country is marked by a large degree of heterogeneity, providing an opportunity to identify the parameters of interest. Indeed, oil price shocks should substantially affect the income distribution in some neighbourhoods, while leaving that of others almost unaffected (at least in the short run). Second, the single payer and universal coverage features of the Canadian health care system facilitate the study of health outcomes at a population-level using administrative data. They make it possible to investigate the dynamics between absolute income, relative income, and health without relying on subgroups that might not yield estimates that are externally valid for other fractions of the population.⁴ Universal health coverage should also limit the probability that (negative) income shocks cause individuals to avoid seeking care for financial reasons, a response which could otherwise bias our estimates by conflating health status and health care needs with access to care. Moreover, while the Canadian health care system differs substantially from the American one in terms of its structure, it does share a number of similarities with other health care systems across the world. In this respect, our results can be used to draw parallels with countries that also have a large social safety net.⁵ Furthermore, to the extent that a correlation between income and access to care may exacerbate the relationships we study, our results can plausibly serve as an informative lower bound for the U.S. context.

Our work contributes to the existing literature on the relative income hypothesis and on the causal dynamics behind the health-income gradient in various ways. First, we focus on the impact of changes to absolute and relative income on more severe health outcomes, such as diagnosed conditions requiring a hospital visit. While most existing work focuses on either mortality, or on measures of subjective well-being and happiness⁶, we focus on health outcomes of complementary

⁴For example, certain studies in the U.S. context must restrict themselves to Medicare or Medicaid populations.

⁵One such example has been documented in the U.K. See Bell and Blanchflower (2018) for a further discussion.

⁶For example on mortality: Gerdtham and Johannesson (2004), Miller and Paxson (2006), Daly et al. (2013), Chetty et al. (2016), Baker et al. (2019), Baker et al. (2021), Case and Deaton (2017). For examples on subjective well-being and happiness: Clark and Oswald (1996), Ferrer-i Carbonell (2005), Luttmer (2005), Becker and Rayo

policy relevance, given their implications in terms of health care utilization and costs, and in terms of individuals' productivity. In addition, objective health outcomes such as the ones we consider do not suffer from most problems of scale and/or comparability commonly associated with happiness and subjective well-being measures⁷, and they allow us to circumvent the pitfalls that come with relying on self-reported health assessments. Second, while most existing work conducted outside of lab settings present correlational relationships between relative income and health, we exploit relatively common oil price shocks to investigate causal dynamics. Instead of relying on cross-sectional variation in average income, or on trends in the mean or median of income distributions at the neighbourhood or region level, our approach focuses on the same individual moving relative positions within a local income distribution. This allows us to gain insights that relate more directly to the health impacts of falling behind, and the mechanisms hypothesized to be behind the *deaths of despair*. Third, our estimating strategy allows us to take into account shifts or changes in the income distribution that might influence health or well-being through other channels than changes in relative income.

Our results suggest that increases in absolute levels of income have a protective effect on individuals' health outcomes in the short-run for men, especially those who are potentially less mobile. We find, however, that changes to a person's relative income have asymmetrical effects on their probability to be hospitalized. For those whose absolute income levels increase following an increase in oil prices, a concurrent increase in their relative income level does not seem to improve health outcomes. On the contrary, our results suggest that for those people, health outcomes could be further improved as more people around them benefit from higher incomes. However, such positive externalities from having richer neighbours are not felt by individuals when their own incomes remain unchanged. Indeed, an increase in hospitalization is observed among those whose relative income decreases following an oil price increase, despite a constant level of absolute income. These dynamics seem to be mostly driven by hospitalizations resulting from drug and alcohol poisonings, mental health and psychiatric conditions – a series of diagnoses consistent with the deaths of despair narrative, and with recent results from the literature on the counter-cyclical movements in mental health distress.

Overall, in addition to providing new insights on the individual-level dynamics linking health trajectories and changes in relative income, the evidence presented sheds new light on the mechanisms through which current trends in inequality and social mobility may cost individuals and societies. The remainder of the paper is organized as follows. Section 2 reviews the relevant literature on associations between health, mortality and well-being and, respectively, absolute and relative income levels. Section 3 investigates the relationship between oil prices and workers labour market outcomes in Canada across industries, while section 4 presents the intuition behind our main empirical approach and section 5 describes the main data sources used to implement it.

(2007), Brodeur and Fleche (2019). Closer to our work is Ifcher et al. (2018), who considers self-reported health outcomes alongside measures of subjective well-being. Our empirical approach, however, differs in that we use an exogenous shock to absolute and relative incomes and consider panel data.

⁷Diamond (2008) summarizes some issues with using subjective well-being or happiness to investigate the relative income hypothesis. Deaton and Stone (2013) also presents an interesting discussion of the particularities of subjective well-being and happiness measures, hedonic or evaluative.

Section 6 introduces our main empirical specification to jointly estimate the health impacts of changes to both absolute and relative income, as well as the associated assumptions that it rests on. Section 7 presents our results, first focusing on a more traditional investigation of the impact of changes to absolute income on individuals' health, and then turning to the separate health effects of changes to absolute and relative income. Finally, section 8 discusses the significance and limitations of our findings, while highlighting avenues for future research.

2 The relationship between income and health

2.1 Absolute income, well-being and health

The relationship between socioeconomic status and physical and mental health has been the focus of a long and rich tradition in economics and beyond. Importantly, the large literature on the social determinants of health documents low income as a predictor of poor health outcomes (Marmot and Wilkinson, 2005; WHO, 2008), and some evidence exploiting natural experiments or policy reforms suggests that a causal channel from income levels to health may be driving at least a portion of this gradient (see Lleras-Muney (2018) for a recent review of important work on the topic).

One of the channels through which low income has been hypothesized to lead to poor health is the allostatic load (McEwen, 1998). Over time, financial challenges and lack of resources can generate hormonal and physiological reactions that weaken the cardiovascular, endocrine and immune systems and contribute to the development of diseases and disorders (Schneiderman et al., 2005).⁸ While the health impact of income shocks can operate through biological channels, the stress they cause could lead to changes in individuals' behaviours (ranging from eating habits (Schneiderman et al., 2005), to increases in smoking and other substance use (Barnes and Smith, 2009), to changes in sleeping patterns and reduced concentration (Forget, 2011)) in ways that might affect their physical and mental health. Tightened or relaxed budget constraints can also, among others, affect individuals' diet quality and diversification in ways that can lead to health changes over time, limit access to health care (especially in the absence of a comprehensive insurance coverage) or the ability to fill prescriptions.⁹

A very rich body of literature has sought to move beyond correlations and to identify the causal impact of own income or wealth on morbidity, health trajectories and even mortality through some or all of these pathways. Studies have also documented the positive short- and long-run effects of cash transfers and tax benefits on the physical and mental health of children (Aizer et al., 2016; Case et al., 2002; Milligan and Stabile, 2009) and mothers (Milligan and Stabile, 2011). Exploiting shocks coming from increases in foreclosure rates Currie and Tekin

⁸Schneiderman et al. (2005) consider exposure to a stressor to be persistent exposure when lasting over a month. Others employ more stringent thresholds.

⁹The literature presents ambiguous evidence on that channel (see, for example, Currie and Tekin (2015)), but we note that financial barriers to health care consumption (excluding medication) are less likely to matter in Canada, given the universal health coverage.

(2015) uncover that decreases in financial well-being and wealth led to non-trivial increases hospital visits and health care utilization for physical and mental health, mostly for emergencies and preventable conditions. Using fluctuations in U.S. stock prices, Schwandt (2018) also finds that decreases in lifetime financial wealth coming from changes in stock prices have negative impacts on retirees' physical health indices, mental health and the probability of living the next two years. However, other work suggests that such higher income may not automatically lead to improvements on measures such as well-being, particularly at the top of the distribution (e.g.: Deaton and Kahneman (2010), Smith (1999)).

Mixed findings, however, emerge from the body of research investigating the impact of changes in economic conditions on physical health. For example, Ruhm (2000) and Strumpf et al. (2017), among others, find evidence of pro-cyclical patterns between mortality, morbidity and macroeconomic trends. Cutler et al. (2006) also present cross-country evidence where income and physical health do not always follow the expected patterns of association.¹⁰ One constant, however, is the observed association (even within the studies cited above) between lower income or economics conditions and the prevalence of mental health distress, suicide, antidepressant consumption, substance-abuse hospitalizations, etc. (for example Kronenberg and Boehnke (2019) in the context of the 2008-2011 financial crisis¹¹, Apouey and Clark (2015) using individual-level data on lottery winners or Bradford and Lastrapes (2014) looking at variations in employment rates).

In addition to this evidence, it is worth mentioning that trends in absolute levels of income do not seem to always match recent trends in mental health or even mortality among population groups. Importantly, Case and Deaton (2015, 2017) suggest that it is the feeling of *falling behind* among less educated non-Hispanic white Americans that could partially be driving the mortality and morbidity trends observed for this group. Such a conclusion, along with evidence that among industrialized countries, differences in mortality patterns are more correlated with the level of inequality within countries than with differences in the average level of income across countries (Strumpf et al., 2017; Baker et al., 2021), suggest that relative income may be an important factor influencing people's health trajectories.

2.2 Relative income, well-being and health

Introduced in Veblen (1899) and later formalized by Duesenberry (1949) as the relative income hypothesis, the idea that relative income might affect individuals' decisions and well-being has a long history.¹² Since then, experimental and clinical studies have documented how the mere

¹⁰Deaton (2012), however, finds strong positive correlations in time series spanning over the 2008 financial crisis, and Ruhm (2015) finds that some of the pro-cyclical relationships between economic conditions and physical health problems observed in the late 20th century are not observed in more recent years.

¹¹Kronenberg and Boehnke (2019) highlight that losses in income are one but do not seem to be the most important factor at play, increased workloads for the employed and the loss of non-wage benefits having stronger impacts.

¹²Duesenberry (1949) suggests that relative incomes can refer to one's income relative to others' (or one's position in the income distribution), but also to one's past experiences and income levels. Our analysis focuses on ideas that relate to the former.

fact of being at the bottom of hierarchies can influence the development of health and/or mood disorders¹³, via changes to the level of stress hormones and HDL cholesterol (Marmot, 1999).

In economics, however, no clear consensus seems to have been reached on the nature of the association between relative income and subjective measures of well-being. On the one hand, some studies have highlighted positive and likely causal relationships between one’s place in the income distribution and measures of self-reported life or work satisfaction (Card et al., 2012; Clark and Oswald, 1996; Clark et al., 2010; Luttmer, 2005), outcomes often linked to mental health. On the other hand, different studies yield more nuanced conclusions. For example, Ifcher et al. (2018) find correlational evidence that the relationship between own well-being and median income within a neighbourhood depends on the level of geography over which the reference group is defined: the association is positive for close neighbours but becomes negative when neighbourhoods are defined at a higher level of aggregation, where effects on public good provision is less likely to dominate.¹⁴ However, they note that these results are reversed for low-income individuals. Moving away from subjective well-being and happiness measures, studies have presented positive associations between relative income and health, longevity, mortality, and suicide (Wilkinson, 1998; Marmot, 2004; Daly and Wilson, 2009; Daly et al., 2013). However, contradictory results on mortality for working-age African Americans have also been presented in the literature (Miller and Paxson, 2006).

In addition to providing mixed results, the empirical literature investigating the contributions of relative income to health and well-being does not, in general, entirely address the potential for reverse causality in the relationship of interest. Moreover, attempts to document the within-individual impacts of moving along the income distribution are fairly limited, with an important fraction of empirical work relying on pooled cross sections of data. Finally, the literature so far has more convincingly argued that long-term exposure to persistent shocks to income, wealth, and other socioeconomic factors can influence health. However, existing work is less conclusive about the short-term health responses to changes in economic circumstances, both absolute and relative, especially in terms of patterns of health care utilization for fully insured populations. Our empirical approach, presented below, seeks to address these outstanding issues.

3 Oil and incomes in Canada

A first objective of our empirical approach is to address the endogeneity of income (both in absolute and relative levels) as an input in the health production function. Recent work has documented the impact of commodity prices on incomes at the individual and aggregate levels.¹⁵ For example, Kilian and Zhou (2018) and Marchand (2015) find that changes in the price of oil have different impacts on employment and incomes among local labour markets, based

¹³Wilkinson (1996) describes this as psychosocial stress.

¹⁴Findings consistent with these are presented in Brodeur and Fleche (2019).

¹⁵An overview of other work linking local labour markets, natural-resource extraction and prices can be found in Marchand and Weber (2018).

on the intensity of local activities related to energy extraction. They note that oil booms have the strongest impact on individuals' labour market outcomes in areas where employment in the oil extraction industry is more prevalent. Fortin and Lemieux (2015) also highlight the role of the energy boom (and of the energy sector) in the wage growth observed in Canadian western provinces in the 2000s compared to provinces less endowed in oil reserves.

In line with these results, our empirical strategy, which we expand on in more detail in the next sections, draws on the hypothesis that changes in oil prices disproportionately affect the labour market outcomes and incomes of individuals working in the oil and gas extraction industry. Since movements in the price of oil are mostly driven by exogenous changes in the global demand and movements in investors' *risk-0* behaviour (Poloz, 2015), we use quarterly variations in the spot price of the West Texas Intermediate (WTI) index as an exogenous shock to income for oil workers. We then rely on these shocks to investigate the health impacts of changes to absolute and relative income.¹⁶ Indeed, while influencing earnings, it is unlikely that oil prices directly influence health outcomes in the short run. This strategy is also in line with the approach presented in Acemoglu et al. (2013), who investigate the impact of income on health care spending at the aggregate level in the U.S. by instrumenting local area income with (lagged) spot oil prices interacted with the intensity of the oil industry at the economic sub-regional level. Numerous other studies have exploited the robust relationship between natural resources prices and income in a variety of contexts. Notably, Brueckner and Gradstein (2016) use global oil price shocks from 1960 - 2007 to estimate the impact of changes in income on schooling.

To investigate how oil prices affect the labour market outcomes of oil workers and non-oil workers in our empirical setting, we turn to the Canadian Labour Force Survey (LFS).¹⁷ The LFS, conducted by Statistics Canada, provides monthly information on workers' labour force status, labour market outcomes and characteristics. It follows respondents for six consecutive months in a rolling panel fashion, selected to form a representative sample of the population age 15 or older. Using individual-level monthly observations from the LFS, we identify individuals whose main labour market activities are in the oil and gas extraction industry, based on the North American Industry Classification System (NAICS). We focus on the years 2006-2010, corresponding to the period covered by our linked census-health data, as described in section 5.

Table 1 presents summary statistics for a representative sample of labour force participants aged 20-60. Compared to their counterparts in other industries, oil workers are disproportionately likely to be male, have on average higher incomes, and tend to work longer hours each week. As shown in table A.1 in the appendix, they are also more likely to work in occupations related to natural and applied sciences (including engineers, science professionals and technicians, site inspectors, regulatory officers, etc.), primary industry occupations (including drillers),

¹⁶Our empirical approach is therefore somewhat related to exposure designs, and while it shares certain similarities with Bartik-style instruments (in reduced form settings), we exploit the local share of oil workers in a different way.

¹⁷In our main empirical analyses, we use the census to document each individual's industry of work. We rely on a cross-section from 2006, and therefore have to turn to another source of data to test the dynamics between oil prices and labour market outcomes.

transport and equipment operators, and other trades (including mechanics, machinists, servicers/maintenance workers). Importantly, only a small proportion (approximately 20% in our pooled cross sections) of those workers are in clerical occupations, or in finance and management jobs. This suggests that most oil workers' incomes are likely to be responsive to fluctuations in the price of oil, as the activity in oil extraction sites picks up or slows down.

Table 1: Characteristics of Workers in the Oil Sector

	Oil Sector	Other Sectors
Male (%)	81.16 (39.10)	50.52 (50.00)
Actual weekly earnings (\$)	1202.10 (853.45)	733.92 (601.96)
Actual hours per week at all jobs	40.07 (22.10)	32.50 (17.22)
Actual hours per week at main job	39.71 (21.90)	31.90 (16.87)
Observations (unweighted)	70,644	2,281,785

Notes: Weighted statistics. Standard deviation in parentheses. Sample of labour force participants aged 20-60, excluding full-time students.

Source: Labour Force Survey 2006-2010, PUMFs, Statistics Canada.

Figures 1 and 2 first provide graphical evidence of the relationship between WTI prices and labour market outcomes across industries, comparing the evolution of unemployment and weekly income for oil workers and for individuals whose professional activities are in other industries.

Figure 1 first suggest that unemployment responds, with a slight lag, to movements in oil prices in a much stronger way for oil workers than for non-oil workers. Figure 2 hints to a similar, although less clear, relationship between oil prices and workers' weekly earnings.

To better understand the dynamics at play, we use the same data in a regression framework. Table 2 presents the impact of the price of oil (with a quarter lag, and scaled by 10) on the same two labour market outcomes, controlling for gender, education, marital status, age as well as province, year and month fixed effects.¹⁸ Results for the full sample are shown in the first

¹⁸Our results are similar (although slightly different in magnitude) when using shorter or longer lags. However, we focus on quarterly lags to match the structure of our main empirical analysis, presented in section 6.

Figure 1: Evolution of the price of oil and unemployment in Canada, by industry





Figure 2: Evolution of the price of oil and average weekly earnings in Canada, by industry



column of panels A and B. They generally confirm the graphical relationships presented in figures 1 and 2: while movements in the WTI index do not seem to influence the weekly earnings and employment status of most workers, they do have a sizable and significant impact on that of individuals employed in the oil and gas extraction sector.¹⁹ In both panels, columns 2 and 3 suggest that the relationships estimated in the full sample is mostly driven by men. This result is not surprising given that the latter are more likely to work in occupations that are more closely tied to the primary sector, and associated with the oil extraction process. This occupational segregation across genders within the oil industry indeed likely contributes to male oil workers' labour market outcomes being more vulnerable to variations in the price of oil.

Table 2: Oil price and labour market outcomes

	Full sample	Men	Women
Panel A: Weekly earnings			
Oil price (one-quarter lag) × Oil industry	51.972*** (19.659)	61.126*** (22.848)	48.544 (34.368)
Oil price (one-quarter lag)	-6.107 (6.793)	-0.852 (10.588)	-10.782 (8.307)
Oil industry	318.000*** (15.273)	310.029*** (17.670)	303.701*** (27.719)
Panel B: Unemployment (marginal effects)			
Oil price (one-quarter lag) × Oil industry	-0.031*** (0.006)	-0.026*** (0.006)	-0.015 (0.013)
Oil price (one-quarter lag)	-0.001 (0.000)	0.004 (0.004)	-0.007* (0.004)
Oil industry	0.014*** (0.005)	0.013** (0.005)	0.003 (0.010)
Observations (unweighted)	2,352,429	1,182,008	1,170,42

Notes: Data for 2006-2010 from the Labour Force Survey, PUMFs, Statistics Canada. Regressions include controls for gender, marital status, age groups, highest educational achievement, province, year and month fixed effects (no CMA-level information available) Price multiplied by 10. Statistical significance: * 10% ** 5% *** 1%

4 Intuition and Framework

Having documented the relationship between oil price movements and incomes in the oil industry, we develop a framework to disentangle the health impacts of a change to one's absolute

¹⁹We note that when looking more specifically at hours worked (not shown), we find that an increase in the price of oil equivalent to \$10 (CAD) increases by nearly two the number of hours worked weekly by oil workers, while not having a similar effect on the time spent at work for individuals active in other industries.

incomes from that of a change in their relative income. An important challenge is that both absolute and relative income tend to move simultaneously. However, given the heterogeneity in the composition of the labour markets across Canadian neighbourhoods, exogenous changes in the price of oil provide us with an opportunity to disentangle the two.

Our identification strategy rests on the following idea: Suppose two individuals, A and B, are employed in the oil industry, but are located in different neighbourhoods. In A's neighbourhood, a majority of people work in sectors of the economy in which earnings do not fluctuate with oil prices (for example education, social services, or public administration). At the other end of the spectrum, B lives in a neighbourhood where the vast majority of individuals also work in the oil industry. As the price of oil rises, both individual A's and individual B's absolute level of income increase. Given that the earnings of most people around her (who are not oil workers) stay the same, individual A's income *relative* to that of her neighbours also increases. However, because B lives in a neighbourhood where the oil industry represents an important share of the labour market, most people around her also see an increase in their income levels. Consequently, and for similar increases in their *absolute* income levels, B's *relative* income does not increase substantially following changes in the price of oil, while A's does.²⁰ We can finally consider the case of a third and fourth individuals, C and D, who are not employed in the oil industry. If C lives in the same (oil intensive) neighbourhood as B, her *relative* income decrease following an increase in the price of oil, despite the fact that she would experience no change in her *absolute* level of income. For D, who lives in the same (non-oil intensive) neighbourhood as A, an change in the price of oil would result in no change to both her absolute and relative incomes.

In summary, our strategy capitalizes on the fact that movements in the price of oil induce different combinations of changes to individuals' absolute and relative incomes, based on their own labour market activity and on the industrial composition of their neighbourhood. Exploiting these combinations, we can separately investigate the impact of changes in absolute and relative income on individuals' health outcomes and utilization of inpatient services.

5 Data

To implement the strategy described above, we use micro data files from the 2006 Canadian Long-Form Census, linked by Statistics Canada to the universe of administrative hospital records from the Canadian Discharge Abstract Database (DAD) for fiscal years 2006 through 2009. Each of these data sets would not, by itself, provide a sufficient basis to investigate the relationship between incomes and health. On the one hand, the Canadian Long-Form Census data provides no information on individuals' health, and it does not enable researchers to follow the same in-

²⁰We note that it is possible that in the long run, earnings in other industries also adjust to changes in the price of oil, consistent with findings in Marchand (2012). However, the figures and results presented in section 3 suggest that in the short run, and for the period covered by our main empirical analysis, this should not be too much of a concern. A discussion of how this could affect our results is nevertheless offered in section 7.

dividuals longitudinally.²¹ On the other hand, while allowing for individuals' health trajectories to be constructed, hospital records do not provide much information on patients' demographic characteristics or on their labour market activities. The combination of those two data sets with information on oil prices coming from the Western Texas Intermediate (WTI) historical price series, however, provides a unique opportunity to shed some new light on the health impacts of changes to absolute and relative income.

5.1 2006 Canadian Long-Form Census

While all Canadians are required to complete a baseline census form every five years, 20% of all respondents are randomly selected to answer an extended set of questions. For this representative sample, information is collected on their demographic characteristics (gender, age, citizenship, highest degree obtained, etc.), their family composition and living arrangements (marital status, family size, renter versus owner, etc.), and their levels and sources of income (for individuals and households) for the year 2006. Importantly in the context of our study, respondents also provide detailed information on their labour market activities, employment status, main occupation and industry of activity. With this last information, we identify all individuals in our sample who work in the oil and gas extraction industry based on three-digit codes from the North American Industry Classification System (NAICS).²²

Crucially for the implementation of our identification strategy, the census files include information on each respondent's geographic location. Following a large portion of the literature investigating the relative income hypothesis using reference groups on the basis of geography, we use this information to assign each individual to a neighbourhood (or local reference group). The combination of respondents' information on their labour market activities and of the detailed geographic markers available for each respondent allows us to create, within each neighbourhood, the share of individuals whose earnings are dependent on the oil and gas extraction industry.

We define an individual's neighbourhood as the census tract of her primary residence. In Canada, census tracts cover census metropolitan areas and census agglomerations²³, and are designed to capture between 2.5 and 8 thousand inhabitants. They are generally drawn according to natural and/or structural boundaries, such that their characteristics can intuitively be interpreted as neighbourhood-level metrics. One caveat of working with census tracts is that they do not completely cover the territory. Indeed, more rural or less densely populated areas are not divided into census tracts. To improve the coverage of our analysis in this context, we treat census agglomerations which are not divided into census tract because of a lower population density to as neighbourhoods as well. As these areas will have a larger surface area and generally include more individuals than census tracts, our empirical analyses include controls for the (log)

²¹Although all individuals are requested to fill in at least a short version of the census questionnaire every five years, a new sample of individuals is randomly selected at each census cycle to answer the long-form questionnaire.

²²We note that our definition of the oil industry does not include individuals working, for example, in refineries, where profitability may be inversely related to the price of oil, the latter being an input in production.

²³Census agglomerations and census metropolitan areas are generally represented as the set of urban areas with a population of at least 50 thousand people.

population in each local reference group in 2006.²⁴ Overall, our main estimating sample should nevertheless be thought of as a primarily urban one.

Results presented in prior work on the relative income hypothesis suggest that the sign of the relationship between others' income and own well-being may depend on the scope or size of the neighbourhoods (or of the reference groups) considered (e.g., Ifcher et al. (2018), Daly et al. (2013), Brodeur and Fleche (2019)). We therefore test the robustness of our results to an alternative (and broader) definition of reference groups, based on census subdivision boundaries. Census subdivisions are defined for statistical reporting purposes, and they generally coincide with the limits of political and administrative geographic units, such as municipalities, cities, towns or villages. While they offer a better coverage of the Canadian territory, including rural areas that are not captured by census tracts, census subdivisions vary substantially in terms of population size and surface area, and are potentially less representative of what individuals' true reference groups may be. They indeed represent more diffuse and less homogenous groupings than census tracts. Using this level of geography could, however, have the advantage of reducing errors when assigning each individual to a reference group. In the coming sections, we therefore compare the results we obtain using both definitions of neighbourhoods.

5.2 2006-2009 hospital records

The 2006 Long-Form Census is merged to all respondents' individual hospital records for the period 2006 to 2009 from the DAD.²⁵ Held and maintained by the Canadian Institute for Health Information, the DAD contains the universe of administrative health records associated with discharges from all inpatient facilities in Canada (except those in the province of Quebec), documenting up to 25 diagnoses flagged for each patient during each stay, and up to 20 health care procedures performed by providers during the patient's hospitalization.²⁶ From these records, we can follow patients through time as they enter and leave inpatient facilities, document their health status at each hospitalization, and across hospitalizations in the case of repeated users. The detailed information on diagnoses allows us to separate diseases, illnesses and conditions responsible for the admission, as well as comorbidities presented by the patient but not responsible for the hospitalization. Using this information, we can look specifically at hospitalizations caused by certain conditions that are more likely caused by financial stress and disruption.

The DAD does not allow us to capture less acute health issues that would lead to an out-patient physician visit rather than an inpatient stay. Our measure of health outcomes therefore only captures the most severe health issues experienced by individuals, which likely impose important costs to individuals and to the health care systems. While our results are thus likely more conservative and may not represent the full health impacts of changes to absolute and relative income, we note that the objective health measures collected in inpatient facilities (hospitals)

²⁴This control, however, becomes irrelevant when we include local reference group fixed effects in our regressions.

²⁵Each year of data starts on April 1, and ends on March 31.

²⁶Diagnoses are identified using the International Classification of Diseases, 10th revision with Canadian enhancements, ICD-10-CA. Procedures are coded using the Canadian Classification of Health Interventions, CCI.

can have a more straightforward interpretation that self-reported health outcomes would.

5.3 Main estimating sample and descriptive statistics

We restrict our main estimating sample to individuals who were aged 20 to 62 in 2006 (as reported in the anchor Census year)²⁷, and for whom labour market and geographic information is available. As mentioned above, we focus on individuals living in census metropolitan areas or census agglomeration such that they can be assigned to a neighbourhood (census tract). We finally transform our data into individual-quarter observations, documenting the number and nature of hospitalizations for each individual in each trimester.²⁸

Because of Statistics Canada data disclosure restrictions, it could be problematic to report descriptive statistics (such as means by subgroups) prior to reaching the stage of submission of the paper for publication. We do, however, note a few important details about our main estimating sample. First, fewer than three percent of individual-trimester observations result in a hospitalization, underscoring the severity of our main health outcome. Second, the population of individuals employed in the oil and gas extraction industry is broadly similar to the general working population with a few notable exceptions. Oil-industry workers in our sample are more likely to be male (as highlighted in section 3), more likely to have a trade education (rather than college or university), and more likely to be employed full-time. While the household size of oil-industry workers is typically (slightly) smaller than the average Canadian household, their income is more than 30% higher in 2006. Given this higher initial average level of income, we expect them to have a better capacity to absorb negative income shocks, at least in the short run, which we keep in mind as we interpret our results.

6 Empirical specification

We develop a reduced-form empirical specification to investigate the dynamics of interest, in a context where individuals' health trajectories can be followed longitudinally but where their labour market outcomes cannot. This feature answers the limitations of the data described in section 5, which prevent us from directly using the price of oil as an instrument for individuals' income in a panel setting. Our main estimating specification, presented in equation 1, rather takes advantage of the information available from the linked census-DAD:

²⁷Our focus on this age sample is motivated by the fact that changes to own income stemming from an oil price shock will affect individuals still likely to receive working income. We therefore exclude any individual who would have reached 65 by the end of our study period. Further, including individuals who are likely to be retirees may introduce noise that could actually introduce a bias in the direction of a relationship between health and own income, as it has been noted that reverse causality from health to income is likely higher for nearly retired or retired individuals (Smith, 1999).

²⁸This choice was made necessary given computing limitations in the research data centre from which the data can be accessed.

$$\begin{aligned}
Hospit_{imt} = & \alpha + \beta_1 \Delta p_t + \beta_2 Oil_i + \beta_3 \Delta p_t \times Oil_i \\
& + \beta_4 \Delta p_t \times Oil_i \times ShareOil_m + \beta_5 \Delta p_t \times ShareOil_m \\
& + \beta_6 ShareOil_m + \mathbf{X}'_{im} \lambda + \mu + \gamma_t + \epsilon_{imt}
\end{aligned} \tag{1}$$

$Hospit_{imt}$ is the health outcome of interest for individual i , living in area m , observed in quarter t . We define this outcome as a binary variable indicating if an individual was hospitalized at least once in quarter t . Oil_i is an indicator equal to one if individual i is employed in the oil and gas extraction industry and $ShareOil_m$ is a measure of the proportion of people in neighbourhood m employed the same industry (and thus more exposed to an oil price shock). Δp_t is the change in the spot price of oil (WTI), which we obtain from the St. Louis Federal Reserve (FRED, 2018). We use the change over a quarter at time t^{29} , and scale the variable to represent a \$10 change in oil prices (given the rare nature of our health outcomes, this facilitates the representation of our results). \mathbf{X}'_{im} includes a set of individual-level controls from the 2006 Census, such as gender, the highest degree obtained, marital status, age (five-year categories), household size, and household income. To capture other time-invariant factors determined at the regional level (such as amenities) and time-varying factors affecting the country as a whole, we also control for CMA/CA fixed effects (μ) and quarter as well as year fixed effects (γ_t).³⁰ We finally also estimate a specification with individual-level fixed effects. In this case, individual-level controls are excluded given that they are measured in a single period during the 2006 Census year.

The relationship between absolute income, relative income and health outcomes is captured by a combination of coefficients in equation 1. First, β_1 corresponds to the impact of a change in the price of oil in the last trimester on the probability that an individual not working in the oil industry is hospitalized. Based on the evidence presented in section 3, we expect this coefficient to be close to zero.³¹ In turn, β_3 captures the differential impact of a change in the price of oil on the health of individuals working in the oil industry. $\beta_1 + \beta_3$ can therefore be interpreted as the health impact of an exogenous shock to absolute income for an individual employed in the oil industry who lives in a local area where everyone else works in a different industry (i.e., where $ShareOil_m = 0$, such that no one else's absolute income gets affected by Δp_t). For an oil worker living in such a neighbourhood, both absolute and relative income would move in a similar way.

Two coefficients capture the relationship between relative income and health. β_4 first speaks to the additional health impact of a change in the price of oil for an oil worker as the share of individuals also working in the oil industry within her neighbourhood increases. Indeed, as the

²⁹We also estimate specifications in which we consider longer lags. Results are consistent with those presented in section 7, and are available upon request.

³⁰We note that our results hold when we replace region fixed effects by province fixed effects, which no longer capture local amenities but still capture important factors such as the features of the health care system, a provincial jurisdiction in Canada. Unlike Daly et al. (2013), who find that the impact of local inequality measures on mortality reverses sign when adding fixed effects for finer geographic units, our results are not affected by the definition given to μ in equation 1.

³¹Despite the fact that they tend to do so substantially less than that of oil workers, the labour market outcomes of non-oil workers could still respond to changes in oil prices, leading β_1 to be different from zero. Evidence presented in section 3 however suggests that this effect should not be too important, all else equal.

value of $ShareOil_m$ (the share of individuals working in the oil industry in the neighbourhood) increases, a disconnect arises between changes in absolute and relative income. In the limit case where $ShareOil_m = 1$, all individuals within neighbourhood m are employed in the oil industry and would see their earnings affected by a change in the price of oil in a similar way. In this case, while each oil worker in the neighbourhood would see her absolute income change, her relative income would remain relatively stable. If relative income does not matter to health, we should expect the value of β_4 to be 0. On the contrary, a positive β_4 and a negative β_3 would suggest the protective effects of an increase in absolute income on health would be weaker when one's relative income does not also increase. Finally, negative values for both β_3 and β_4 would suggest positive health externalities for those surrounded by individuals whose absolute income increases alongside theirs.

Still in equation 1, β_5 encapsulates the health impact of a change in the price of oil for non-oil workers as the share of their neighbours employed in the oil industry (and whose incomes thus vary with Δp_t) increases. For these non-oil workers, whose absolute income remains constant, a change in relative income of the opposite sign as the oil price movements becomes increasingly important as $ShareOil_m$ gets closer to 1. Accordingly, a positive value of β_5 would suggest that a decrease in one's relative income increases their probability of hospitalization. On the other hand, a negative value of the coefficient would suggest that being surrounded by people whose income increases, even when one's own income doesn't, can generate positive health externalities.

A perhaps more insightful way to think about the results from equation 1 is by re-estimating the model while centering the variable $ShareOil_m$ around different values Z (between 0 and 1). We can then compare the evolution of β_1 and β_3 across estimations. The change in the estimated value of β_3 obtained when $ShareOil_m$ is re-centred around higher values Z allows us to understand how the health impact of an exogenous shock to absolute income for an oil worker is strengthened/attenuated as the coincidental shock to relative income dissipates. A similar reasoning can be applied to the variations in the estimated values of β_1 as $ShareOil_m$ gets re-centred around higher values. Increasing values of the estimated β_1 as $ShareOil$ is centred around higher Z values would suggest that decreases in one's relative income can have adverse health effects (increasing their probability of hospitalization).³²

It is worth nothing that, since we do not have longitudinal labour market outcomes and demographic information for the individuals in our dataset, four assumptions are important for the identification of the parameters in our reduced form approach. First, it is key for the identification of our model that changes in oil prices disproportionately affect the incomes of oil workers compared to those of individuals employed in other industries. Evidence from the Labour Force Survey presented in section 3 provides evidence that it is the case. Second, the industrial composition of the neighbourhoods in our sample should remain relatively stable between 2006 and 2010. Third, individuals in our sample generally don't move neighbourhoods during the period covered in our analysis. Finally, oil prices do not affect individuals health in other ways than

³²While generating information similar to that conveyed by β_4 and β_5 but in a more intuitive way given the triple interaction with continuous terms, this approach also gives a better idea of the value of $ShareOil_m$ that needs to be reached for relative income to matter (if at all).

through changes to individuals’ absolute or relative incomes. We further discuss potential threats to identification linked to these assumption in section 7.

Our empirical approach differs from most empirical work investigating the health implications of the relative income hypothesis. The literature generally includes a linear term in the first moment of the income distribution of reference (or in the median income) as a control in regression equations. A negative coefficient on this variable is often interpreted as evidence of lower relative income impairing one’s health or well-being (in line with Wilkinson (1996)’s psychosocial stress theory). One challenge with such specifications, however, is linked to the identification of potential complementarities between changes in individuals’ absolute and relative incomes, and the differentiation between the opposing forces that positive externalities from higher neighbourhood income (through better amenities, etc.) and lower relative income may represent. Equation 1 seeks to address these challenges, while focusing on the health impacts of individual-level *changes* to absolute and relative income.

7 Results

7.1 The health impacts of absolute income changes

To get a better picture of the dynamics at play, we start by estimating a simplified version of equation 1 which ignores the potential health effects of changes to relative income. This helps to narrowing in on the relationship between a change in absolute income (coming from changes in oil prices) and individuals’ probability of hospitalization. Apart from the absence of terms referring to the variable $ShareOil_m$ in equation 1, all variables in equation 2 have the definition previously given.

$$Hospit_{imt} = \alpha + \kappa_1 \Delta p_t + \kappa_2 Oil_i + \kappa_3 \Delta p_t \times Oil_i + \mathbf{X}'_{im} \lambda + \mu + \gamma + \epsilon_{imt} \quad (2)$$

The estimated coefficient κ_1 conveys information on the effect of a change in the price of oil on the probability of hospitalization for a non-oil worker, whose absolute level of income should not be impacted as highlighted in section 3. Accordingly, we expect estimated values of κ_1 to be relatively close to 0. However, we would expect κ_3 , the impact of a change in the price of oil on the probability that an oil worker gets hospitalized, to be negative if increases in absolute levels of income have a protective effect on individuals’ health.

The results from a linear probability estimation of equation 2 are presented in table 3. The first two columns focus on the full sample of individuals in the data, and display the expected coefficients: while a change in oil prices does not seem to influence individuals’ likelihood of being hospitalized when they do not work in the oil industry ($\hat{\kappa}_1 \approx 0$), it is inversely related to the probability of inpatient stays for oil workers ($\hat{\kappa}_3 < 0$ and statistically significant). Given that oil prices exogenously shock the absolute levels of income for oil workers, we interpret these results as evidence of a small but causal impact of own-income dynamics on health. According to our estimates, an increase of \$100 in the price of oil int the last trimester decreases the probability

that an oil worker gets hospitalized by 0.1%, a reduction of close to five percent in the baseline probability of hospitalization in any given quarter in our data. We note that the coefficients on other individual-level control variables have the expected signs. For instance, having a regular salaried job is negatively associated with the probability of hospitalization, and younger age groups are less likely to be hospitalized.³³ The results are robust to including individual fixed effects (column 2).

Table 3: Probability of Hospitalization (absolute income analysis)

	Full sample		Men only	
Δp_t (κ_1)	0.0000 (0.0000)	0.0000 (0.0000)	0.0003 (0.0002)	0.0003 (0.0002)
Oil_i (κ_2)	-0.0010*** (0.0003)		-0.0005* (0.003)	
$\Delta p_t \times Oil_i$ (κ_3)	-0.0010** (0.0004)	-0.0010** (0.0004)	-0.0022** (0.0009)	-0.0022** (0.0010)
CMA/CA Fixed Effects	✓		✓	
Individual Fixed Effects		✓		✓
Observations	23,310,670	23,310,670	11,334,360	11,334,360

Notes: Statistical significance: * 10% ** 5% *** 1%. Standard errors are clustered at the CMA/CA level. All columns include controls for immigrant status, single parent, marital status, number of individuals in the economic family (log), income and income² (log), as well as CMA/CA, quarter and year fixed effects. Δp_t is scaled up by 100.

As noted before, our estimates in table 2 reveal that the relationship between income and oil prices is particularly pronounced for men, and quite imperceptible for women. Given this heterogeneity, we present in columns 3 and 4 of table 3 the results from the estimation of equation 2 on a men-only sample. The estimated coefficients for this exercise tell a similar story, although the relationship between own income and health, captured by κ_3 , are more than twice as strong. Here again, the results are robust when including individual fixed effects in the specification.

While interesting, and in line with previous work suggesting a causal pathway from absolute income levels and individuals' health, these results do not inform on – and might actually be hiding – the coincidental health effects of changes to individuals' relative income, an issue we

³³For detailed results, see the corresponding table A.2 in the appendix.

address next.

7.2 The Health Impacts of Relative and Absolute Income Changes

We then turn to our main empirical specification, which considers the impact of both changes to absolute and relative income on the probability of hospitalization. Given the results presented in section 3 and 7.1, we focus on the sample of working age men to study these dynamics.³⁴ As mentioned in section 6, our preferred approach is to estimate a version of equation 1 that includes CMA/CA fixed effects, and with the variable $ShareOil_m$ centered around different values $Z \in [0, 1]$. To choose the values Z to use, we first sort all neighbourhoods from the one with the smallest to the one with the largest share of oil workers. We then select the neighbourhoods corresponding to several points of the distribution (deciles, 95th, 99th and 99.9th, the latter corresponding to a 24% share of oil workers in the neighbourhood). Finally, we retrieve the values Z corresponding to the share of oil workers in these neighbourhoods.³⁵


Across estimations, we then start by studying the evolution of the coefficient β_3 as we move from lower to higher levels of Z . On the far left of figure 3, the coefficient represent the impact of a change in oil price on the health of an oil worker living in a neighbourhood populated predominantly by non-oil workers ($Z = 0$, such that the individual experiences a change in both absolute and relative income). As we move towards the right, the coefficient represents the health impact of a change in oil price for an oil worker who lives in a neighbourhood with a larger share of oil workers (such that we should expect his absolute but not his relative income to vary). Although they are small and generally not statistically significant, all estimated β_3 reported on the graph are negative. This is consistent with the dynamics mapped in figures 1 and 2, and with the robust relationship between earnings and oil prices identified in 2: as increases in oil prices have a positive impact on labour market outcomes for oil workers, the associated increase in income could have mild protective effects on oil workers' health, and reduce their probability of hospitalization. Moving from the left to the right of the figure, however, provides some interesting and perhaps unexpected insights: the coefficient on $\Delta p_t \times Oil_i$ (β_3) remains negative and nearly triples in magnitude for neighbourhoods with a larger share of oil workers. In the top decile of most oil-intensive neighbourhoods, a \$10 increase in the price of oil in the last trimester could reduce the probability of hospitalization for an oil worker by 0.7% to 2.3%. In other words, for those whose absolute incomes move with changes in the price of oil, being surrounded by people experiencing similar increases (decreases) in income generates some positive (negative) externalities on their health outcomes. Increases in relative income as one's absolute income increases would therefore not yield additional health benefits.³⁶

³⁴The results for the full sample of men and women are presented in the Appendix. While they follow similar patterns as the ones obtained on a sample of men, they tend to be weaker and indistinguishable from zero.

³⁵The variation in the share of oil workers at the lower end of the distribution of neighbourhoods is quite minimal. For example, the bottom 20 percent of census tracts have extremely low fractions (close to 0) of individuals employed in the oil industry. The local shares, however, increase much more quickly at the top of the distribution of neighbourhoods.

³⁶This result corresponds to a negative estimated value of β_4 in equation 1. Results available upon request.

Figure 3: Oil workers: Impact of changes in oil prices on the probability of hospitalization (β_3 for all men, across neighbourhoods with increasing concentration in oil workers)




Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, marital status, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.

We then conduct a similar exercise with β_1 . On the far left of figure 4, the estimated coefficient represents the health impact of a change in oil price on the health of a non-oil worker living in a neighbourhood populated predominantly by non-oil workers ($Z = 0$, such that no change in absolute or relative income should be experienced by the individual). Although slightly positive, this impact is very small and not statistically significant. This result is in line with the assumption (and the evidence from section 3) that a change in the price of oil should not affect the incomes of non-oil workers in the first place. However, as we move to the far right of the figure, the estimated coefficient can be interpreted as the health impact of the same change in oil prices for a non-oil workers who lives in a neighbourhood with a significant share of oil workers (such that, as the change in oil price increases, the absolute income of the non-oil worker remains the same, and his relative income decreases). Although all plotted coefficients remain indistinguishable from zero, the movements in β_1 along the x-axis is consistent with a situation where for individuals whose absolute income remains unchanged, a decline in relative income (or the feeling of *falling behind*), may be harmful for their health (and vice versa).

While figures 3 and 4 present interesting patterns, most of the findings are not precise enough to be informative. However, one of the assumptions listed at the end of section 6 is less likely to be satisfied among certain demographics, which could bias our estimates of β_1 toward zero. Indeed, it is possible that non-oil workers would choose to either move away from an oil-intensive neighbourhood as the price of oil increases, to avoid the feeling of *falling behind*.³⁷ As such relo-

³⁷Green (2010) and Clark, Frijters and Shields (2008) do provide evidence that mobility may be inversely related

Figure 4: Non-oil workers: Impact of changes in oil prices on the probability of hospitalization (β_1 for all men, across neighbourhoods with increasing concentration in oil workers)



Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, marital status, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.


cations to less oil-intensive neighbourhoods should be easier for single men, who would not have to move a full family with them, we next focus our analysis specifically on married men. In our sample, single men are indeed three times more likely to report having moved in the last three years than married men. We thus present the equivalent of figures 3 and 4, for married men only. In addition, the repercussions of absolute and relative income changes on single individuals may be more muted, since the latter are less likely to fear their repercussions on other family members.

Figure 5 first considers married male oil workers. The dynamics shown are quite similar to those presented for all men, although at times nearly twice as large in magnitude. While the impact of a change in the price of oil in neighbourhoods with few oil workers is not statistically different from zero, a \$10 increase in the price of oil for an oil worker living surrounded by a critical mass of individuals experiencing similar positive income shock decreases their probability of being hospitalized by 1.3%-4.3% . Here again, it seems that the positive externalities from being surrounded by richer people are felt among those whose own income have also increased. However, the situation is reversed for those who are not lucky enough to see their own income increase. For them, figure 6 suggests that, in the absence of an increase in absolute income, a decrease in relative income coming from a \$10 increase in oil prices could lead to a 0.4%-1.6% increase in the probability of hospitalization. Although a change in the price of oil does not in and of itself affect the probability that a non-oil worker gets hospitalized (as it does not affect her level of income in absolute terms), it can make him worse if he lives in a neighbourhood with

to feeling of economic and job satisfaction.

a sufficiently high proportion (here, 20-24%) of oil workers. We note that the effects we estimate are still quite small, and only different from zero in the few neighbourhoods heavily populated by individuals employed in the oil industry. However, we note that these results need to be taken in the context of a very low baseline probability of hospitalization in our data. Indeed, in a given quarter, fewer than three percent of our census sample has a record for an inpatient stay.

Figure 5: Oil workers: Impact of changes in oil prices on the probability of hospitalization (β_3 estimates for married men, across local concentration of oil workers)



Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.


Perhaps unsurprisingly given the nature of our data and the type of variations we exploit (exogenous changes in the price of oil), our results remain similar when we exploit the longitudinal nature of the DAD to estimate a versions of equation 1 with individual fixed effects.

7.3 Sensitivity analyses

7.3.1 Alternative reference groups and functional forms

As mentioned in section 6, previous work has shown that the relationship between relative income and health can be sensitive to the choice of reference group, especially when they are defined over geographic areas. To test if our own conclusions are sensitive to the nature and size of the neighbourhoods we consider, we re-estimate equation 1 using census subdivisions rather than census tracts, the former being substantially larger and more populous (and exhibiting more heterogeneity in population size). Our results, both with CMA/CA fixed effects and with individual fixed effects, are qualitatively similar when opting for this choice of geography, although they are slightly smaller in magnitude and less precisely identified. This may not come as a

Figure 6: Non-oil workers: Impact of changes in oil prices on the probability of hospitalization (β_1 estimates for married men, across local concentration of oil workers)



Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.

surprise, as the entire city core of Toronto (approximately 2.5 million people) corresponds to a single census subdivision, potentially not representing accurately individuals' reference group.

We also note that while our main results are estimated using a linear probability model, focusing on the probability that an individual gets hospitalized at least once in a trimester, they are robust to using a probit specification. Moreover, we can investigate the severity of deteriorations in individuals' health by looking at the number of hospitalizations, rather than on the probability of being hospitalized. Figures A.3 and A.4 in the appendix present the results from the estimation of an inverse hyperbolic sine transformation model that allows us to simultaneously treat a count dependent variable with a mass at zero. Very few individuals being hospitalized more than once in a given time period, the results are very similar to those obtained with the linear probability model.

7.3.2 Changes in purchasing power as an alternative mechanisms

So far, we have interpreted our results on β_1 as the impact of changes to relative income of married men's health outcomes. However, an alternative mechanism could be consistent with these results, and cannot directly be ruled out by our empirical approach. It is possible that as the general level of income increases in oil-intensive neighbourhoods in response to rising oil prices, local prices in these neighbourhoods adjust accordingly. For non-oil workers living in these areas, whose own incomes have not increased, this local inflation could then result in a

Table 4: Oil prices and local inflation
(Oil-intensive and non oil-intensive Canadian cities)

	CPI all items		CPI shelter	
Oil price (one quarter lag)	0.020*	0.017*	0.040***	0.040***
	(0.011)	(0.008)	(0.009)	(0.009)
Oil price(one quarter lag) \times Oil intensive		0.004		0.014
		(0.018)		(0.050)
Observations	270	270	270	270

Notes: Regressions include year, month and city fixed effects. Monthly observations. 2006-2011. Statistical significance: * 10% ** 5% *** 1%

reduction in their purchasing power. With this type of general equilibrium dynamics, our results on β_1 could reflect the health impacts of such a decline in the capacity to maintain a certain level of consumption, rather than the health impacts of a decrease in relative income.

One way to test for this alternative mechanism is to investigate how different prices indices (such as rent, the local consumer price index, etc.) evolve alongside our WTI index, across local markets and throughout our sample period. While we do not have access to census-tract level price data, we gather city-level monthly price indices from Statistics Canada for some of the larger cities included in our main estimating sample. Figure 7 maps the evolution of the average local all-items consumer price index separately for oil-intensive (Edmonton and Calgary) and non-oil intensive (Montréal, Québec, Vancouver, Winnipeg, Ottawa-Gatineau, Toronto and Hamilton) geographical areas between 2005 and 2013. While price levels are systematically higher in non-oil intensive areas (reflecting the reality of larger cities in that group, such as Toronto and Vancouver), they seem to follow a trajectory that is quite similar to that of price levels in oil-intensive cities. The first column of table 4, which presents the corresponding regression results with the lag structure used in our main empirical models, confirms this absence of differential relationship between oil prices and local price levels between oil-intensive and other Canadian cities. In fact, lagged oil prices seem to be associated with higher local prices overall, but not particularly so where oil workers are concentrated. Figure 8 and the second column of table 4 reproduce similar results for local shelter costs, one of the most important expenditures made by households, and a budget item that could directly affect individuals' purchasing power. Although imperfect, this evidence lends credibility to our interpretation of our results from equation 1.

It is also important to mention that our analysis focuses on a period that coincides with a recession, which could also have its own impacts on individuals' health outcomes. However, we note that the impact of the economic crisis would be felt among vast swaths (if not the

Figure 7: Movements in oil prices and in consumer price index
(Oil-intensive and non-oil intensive Canadian cities)

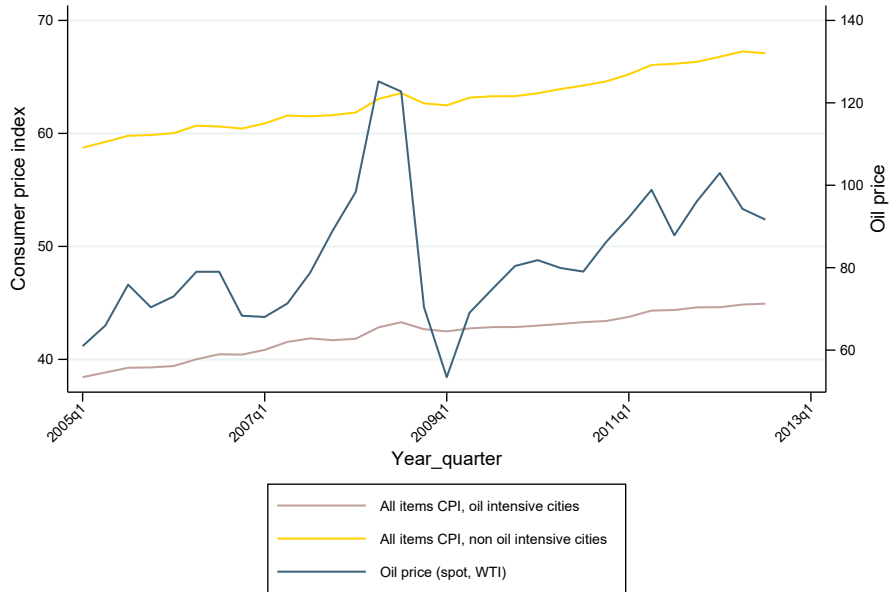



Figure 8: Movements in oil prices and in shelter costs
(Oil-intensive and non-oil intensive Canadian cities)



majority) of the population, and across industries. Since our main parameters of interest compare individuals based on their industries of work and on the concentration of oil workers in their neighbourhoods, while controlling for time fixed effects, this coincidental event should not interfere with our estimation strategy.

7.3.3 Falsification tests

One alternative way to test the interpretation of our results is to reproduce our empirical exercise while using placebo industries in lieu of the oil industry. First, we check that a few other specific industries' incomes are not correlated with oil prices. Nursing and residential care facilities and hospitals are an interesting example. First, they are characterized by incomes that display no co-movement with the WTI. Second, these industries also employ a share of Canadian workers that is quite similar to that of the oil industry and although they don't present as strong geographical concentration patterns, there is a variability in the share of workers they employ across neighbourhoods.

Figures A.5 and A.6 in the appendix show the evolution of β_1 and β_3 on a sample of married men when we replace the variables Oil_i and $ShareOil_m$ in equation 1 to convey similar information but related to workers ties' to the Nursing and residential care facilities and hospitals industries. The gradient that could be observed in our main results is no longer perceptible here, and if anything, the patterns observed (although all not statistically significant) go in the opposite direction, suggesting that the assumptions behind and mechanisms implied by our empirical framework seem credible. Reproducing this exercise with the non-professional, scientific and technical services industries yields similar results (figures A.7 and A.8 in the appendix). This time, the concentration of all estimated values of β_3 and β_1 around zero and no visible gradient as we move across neighbourhoods with different concentration of workers in these industries.

7.4 Potential threats to identification

As highlighted in section 6, our identification strategy relies on four assumptions, which come from the nature of our data. First, it assumes that a change in the price of oil (Δp_t) will disproportionately affect the income of individuals working in the oil and gas extraction industry, a relationship for which we provided supportive evidence from a representative sample in section 3. However, we do not completely rule out the possibility that increases in oil prices may lead to positive spillovers on earnings in other industries. Importantly, we note that Fortin and Lemieux (2015) suggest that positive wage spillovers from oil booms are more likely to arise in neighbourhoods or local labour markets where the oil industry was initially more prevalent.³⁸ In the context of equation 1, this would imply that in response to an increase in oil prices, incomes should increase for a larger portion of the population than what we assume in neighbourhoods where the oil industry represents an important share of the labour market. If the magnitude of spillovers within a local labour market is monotonic in the proportion of oil workers, these dynamics should not contradict our interpretation of the results. Under such circumstances, it is

³⁸Similar results are also presented for oil booms and busts in the United States (e.g., Allcott and Keniston (2018)).

also possible that we incorrectly identify certain non-oil workers as having their absolute income *not* affected by changes in oil price, which would work against us estimating an effect close to 0 when considering neighbourhoods with a low concentration of oil workers.

Second, our main estimating equation assumes no change in the industrial composition of neighbourhoods during our study period. In contrast, Fortin and Lemieux (2015) note a growth in employment in the Canadian oil and gas industry as the price of oil was surging in the 2000s, leading local labour markets that were already oil-intensive at the beginning of this period to see a further increase in their industrial concentration. Our data does not allow us to track changes in the share of oil workers by census tracts beyond 2006. However, we note that in the changes described by Fortin and Lemieux (2015), employment shares in the oil industry disproportionately increase in areas where oil already played a predominant role in the labour market, leading to dynamics similar to the ones addressed in the previous paragraph.³⁹

One more important caveat of our analysis, related to a third identifying assumption of our approach, is that we cannot directly address the potential for individuals to move across neighbourhoods after 2006, which would lead to some measurement error in our empirical exercise and to potential biases if those moves are not random. In the latter case, and in line with prior findings on workers' mobility, we expect that non-oil workers in non-oil-intensive neighbourhoods would move to oil-intensive neighbourhoods to become oil workers when the price of oil surges. A first consequence of such dynamics would be to work against our findings, as we would assign individuals to the wrong industries. A second consequence would be the underestimation of oil workers in oil-intensive neighbourhoods as the price of oil increases, an issue that once again is quite similar to the one discussed in the previous paragraph. We, however, note that most individuals with repeated inpatient stays in our data visit the same facilities, and are observed within the same province, suggesting that at least inter-city and inter-provincial mobility should be relatively limited in our sample. Additionally, our results are robust to exclusively looking at a sample of men who reported not having moved in the five year prior to the 2006 census, which is reassuring if we think that past moves are predictive of the propensity to move after 2006.

A final identifying assumption is that an increase in oil prices does not affect people's health through other channels than their absolute or relative income (the exclusion restriction). This assumption could be violated, for example if an increase in the price of oil caused changes in air quality at the local level that had an impact on the frequency and volume of hospital visits for chronic obstructive pulmonary diseases, asthma or other chronic conditions, independently from changes in individuals' income levels. McLinden et al. (2012) indeed present evidence of greater volume and concentration of NO_2 (and potentially SO_2) over and around areas corresponding to larger oil sands mining and extraction sites in Canada between 2005 and 2010, for which the activity level is positively correlated with the WTI index. The U.S. Environmental Protection Agency warns against the short- and long-run adverse effects that these two types of pollutants can have on the aggravation of respiratory diseases, with a potential to lead to

³⁹We further note that Kilian and Zhou (2018) find limited evidence of mobility from oil-poor to oil-intensive regions in Canada following a surge in oil prices, a finding echoing similar patterns found in the context of trade shocks (e.g., Dix-Carneiro and Kovak (2017)).

hospital admissions and emergency room visits (EPA, 2018a,b). Schlenker and Walker (2016) also present evidence that punctual increases in CO exposure due to excess airplane idling on tarmacs have short-term impact on emergency room visits and hospitalizations within neighbouring communities. Although these dynamics are important to consider, we note that the exclusion restriction should mostly be satisfied when considering an outcome variable that excludes hospitalizations related to respiratory problems, which we do in a series of robustness checks. It is also likely that these effects of oil prices on health through increased oil extraction activities are expected to be mostly concentrated among vulnerable populations such as children and seniors (Schlenker and Walker, 2016), groups that are excluded from our main estimating sample.⁴⁰

7.5 Potential pathways and mental health

Our results so far hint at the fact that, when an individual witnesses her neighbours' income increase but their own doesn't, the feeling of falling behind may have a slight negative impact on their health. This is consistent with the intuition presented in Case and Deaton (2015), according to which the feeling of falling behind might be related to rising morbidity and mortality among certain groups in the United States. To investigate this possibility further, we focus our analysis on a select set of hospitalizations caused by factors associated with the concept of the deaths of despair: drug and/or alcohol poisonings, admissions for which alcohol is the primary cause, self-harm, as well as admissions for mental health or psychiatric conditions. Since the volume of hospitalizations for those specific motives are quite rare, we slightly alter our empirical strategy to explore the probability that a given hospitalization is caused by one of the identified *conditions of despair*. In the light of the results presented above, we continue to focus on the subsample of married working-age men.

The results from these estimations⁴¹ highlight an interesting pattern: when moving from a reference neighbourhood with a low share of oil workers to one with a high share of oil workers, the effect of an oil price shock on the probability that a given hospitalization is caused by a condition of despair goes from being negative and statistically significant to being positive and statistically significant for non-oil workers. In other words, conditional on hospitalization, the probability that a non-oil worker is admitted to an inpatient facility for a reason related to drugs, alcohol, self-harm or mental health goes from being significantly lower when her relative income does not change to being significantly higher when she lives in an environment where a fifth of her neighbours see their incomes increase while she doesn't (such that her relative income decrease).

For oil workers, an increase in the price of oil (and thus in their absolute income) seems to reduce the probability that a given hospitalization is caused by any of the conditions listed above. However, the estimated coefficients are small and not statistically significant, preventing

⁴⁰Schlenker and Walker (2016) also find weaker relationship between air pollution exposure and inpatient stays, compared to other health outcomes and measures of health services utilization.

⁴¹These results are in the process of being vetted by Statistics Canada for release from the research data centre, and should be available in the next iteration of the draft.

us from making strong inference to that effect.

8 Conclusion

Despite policies designed to foster redistribution, the concentration of income among top earners has increased in most developed economies, as has inequality within and across communities. Aside from the well-documented repercussions of these trends on social cohesion and productivity, it has been suggested that changes in the income distribution might have an impact on individuals' physical and mental health. For example, recent work has pointed to the feeling of *falling behind* as a potential driver for the rise in morbidity and mortality among certain groups of Americans, for whom absolute levels of income or absolute economic conditions have not necessarily declined, but whose outcomes relative to that of others may have.

In this paper, we propose an empirical strategy that draws on the importance and geographic concentration of the oil extraction industry in Canada to identify the health effects of changes to people's relative economic situation from changes in their absolute level of income. First, to deal with the fact that both absolute and relative incomes are likely endogenous inputs in the health production function, we exploit exogenous movements in the price of oil, which predominantly affect the earnings of individuals employed in the oil industry. These variations in the price of oil also induce movements in local income distributions, based on the makeup of local labour markets and on the share of individuals employed by the oil industry. This allows us to deal with a second identification challenge: changes in individuals' income levels often simultaneously affect their position within the income distribution. Our strategy capitalizes on the fact that movements in the price of oil induce different combinations of changes to individuals' absolute and relative incomes, based on their own labour market activity and on the industrial composition of their neighbourhood. Exploiting these combinations allows us to shed some light on the extent to which changes in absolute and relative income trajectories separately contribute to the development of severe health conditions and the use of inpatient care services.

Our results suggest that increases in absolute income levels coming from increases in oil prices can have small protective (although not always statistically significant) effects on individuals' health, as measured by a reduction in the probability of hospitalization. They also reveal interesting asymmetrical effects of changes to relative income across groups, especially among married men. We find that for those who experience an increase in their absolute levels of income, being surrounded by people whose economic situation also improves is accompanied by positive health externalities, which can lead to statistically significant reductions in the probability of hospitalization. However, a slight increase in the probability of hospitalization can be observed for individuals whose absolute income doesn't change following a movement in the price of oil, but who experience a drop in relative income because of increases in their neighbours' income levels. Hospitalizations for drugs or alcohol abuses, liver diseases and mental health problems seem to be important contributors to these dynamics. We note that these results are qualitatively similar but not statistically significant when focusing on all working age men, and do not seem to apply

to women. This last observation is coherent with evidence according to which male oil workers' incomes respond with more sensitivity to changes in oil prices, and to the lower mobility of married men in responses to changes in the economic conditions within their reference group.

Our results shed new light on the mechanisms through which movements in individuals' income, and more broadly income inequality may affect mental health care use and growing disease burdens. These findings point to important channels through which the deaths of despair phenomenon may be operating. Our empirical approach is, however, characterized by limitations. Importantly, our data does not allow us to investigate long-term dynamics, general equilibrium effects, or to explore the extent to which specific channels (social networks, allostatic load, changes in health behaviours such as diet, exercise, etc.) may lead income changes to alter individuals' health outcomes. These offer promising avenues for future work.

References

- Acemoglu, D., A. Finkelstein, and M. J. Notowidigdo (2013). Income and health spending: Evidence from oil price shocks. *Review of Economics and Statistics* 95(4), 1079–1095.
- Aizer, A., S. Eli, J. Ferrie, and A. Lleras-Muney (2016). The long-run impact of cash transfers to poor families. *American Economic Review* 106(4), 935–971.
- Allcott, H. and D. Keniston (2018). Dutch disease or agglomeration? The local economic effects of natural resource booms in modern america. *Review of Economic Studies* 85(2), 695–731.
- Apouey, B. and A. E. Clark (2015). Winning big but feeling no better? the effect of lottery prizes on physical and mental health. *Health Economics* 24(5), 516–538.
- Baker, M., J. Currie, B. Miloucheva, H. Schwandt, and J. Thuilliez (2021). Inequality in mortality: updated estimates for the united states, canada and france. *Fiscal Studies* 42(1), 25–46.
- Baker, M., J. Currie, and H. Schwandt (2019). Mortality inequality in Canada and the United States: Divergent or convergent trends? *Journal of Labor Economics* 37(S2), S325–S353.
- Barnes, M. G. and T. G. Smith (2009). Tobacco use as response to economic insecurity: Evidence from the National Longitudinal Survey of Youth. *B.E. Journal of Economic Analysis and Policy* 9(1), 1–29.
- Becker, G. S. and L. Rayo (2007). Evolutionary efficiency and happiness. *Journal of Political Economy* 115(2), 302–327.
- Bell, D. N. and D. G. Blanchflower (2018). The well-being of the overemployed and the underemployed and the rise in depression in the UK. Technical report, National Bureau of Economic Research.
- Bradford, W. D. and W. Lastrapes (2014). A prescription for unemployment? Recessions and the demand for mental health drugs. *Health Economics* 23, 1301–1325.
- Brodeur, A. and S. Fleche (2019). Neighbors’ income, public goods, and well-being. *Review of Income and Wealth* 65(2), 217–238.
- Brueckner, M. and M. Gradstein (2016). Income and schooling: Evidence from international oil price shocks. *Journal of Human Capital* 10(2), 212–234.
- Card, D., A. Mas, E. Moretti, and E. Saez (2012). Inequality at work: The effect of peer salaries on job satisfaction. *American Economic Review* 102(6), 2981–3003.
- Case, A. and A. Deaton (2015). Rising morbidity and mortality in midlife among white non-hispanic americans in the 21st century. *Proceedings of the National Academy of Sciences* 112(49), 15078–15083.
- Case, A. and A. Deaton (2017). Mortality and morbidity in the 21st century. *Brookings Papers on Economic Activity Spring 2017*, 397–476.

- Case, A., D. Lubotsky, and C. Paxson (2002). Economic status and health in childhood: The origins of the gradient. *American Economic Review* 92(5), 1308–1334.
- Chetty, R., M. Stepner, S. Abraham, S. Lin, B. Scuderi, N. Turner, A. Bergeron, and D. Cutler (2016). The association between income and life expectancy in the United States, 2001-2014. *JAMA* 315(16), 1750–1766.
- Clark, A. E., N. Kristensen, and N. Westergaard-Nielsen (2010). Economic satisfaction and income rank in small neighbourhoods. *Journal of the European Economic Association* 7(2-3), 519–527.
- Clark, A. E. and A. J. Oswald (1996). Satisfaction and comparison income. *Journal of public economics* 61(3), 359–381.
- Currie, J. and E. Tekin (2015). Is there a link between foreclosures and health? *American Economic Journal: Economic Policy* 7(1), 63–94.
- Cutler, D., A. Deaton, and A. Lleras-Muney (2006). The determinants of mortality. *Journal of Economic Perspectives* 20(3), 97–120.
- Daly, M. C. and D. J. Wilson (2009). Happiness, unhappiness, and suicide: An empirical assessment. *Journal of the European Economic Association* 7(2-3), 539–549.
- Daly, M. C., D. J. Wilson, and N. J. Johnson (2013). Relative status and well-being: Evidence from US suicide deaths. *Review of Economics and Statistics* 95(5), 1480–1500.
- Deaton, A. (2012). The financial crisis and the well-being of Americans. *Oxford Economic Papers* 64, 1–26.
- Deaton, A. and D. Kahneman (2010). High income improves evaluation of life but not emotional well-being. *Proceedings of the National Academy of Sciences of the United States of America* 107(38), 16489–16493.
- Deaton, A. and A. A. Stone (2013). Two happiness puzzles. *American Economic Review: Papers and Proceedings* 103(3), 591–597.
- Diamond, P. (2008). Behavioural economics. *Journal of Public Economics* 92, 1858–1862.
- Dix-Carneiro, R. and B. K. Kovak (2017). Trade liberalization and regional dynamics. *American Economic Review* 107(10), 2908–2946.
- Duesenberry, J. S. (1949). *Income, saving, and the theory of consumer behavior*. Cambridge (MA): Harvard University Press.
- EPA (2018 (accessed December 14,2018)a). *Nitrogen Dioxide (NO₂) Pollution*. United States Environmental Protection Agency. <https://www.epa.gov/no2-pollution>.
- EPA (2018 (accessed December 14,2018)b). *Sulfure Dioxide (NO₂) Pollution*. United States Environmental Protection Agency. <https://www.epa.gov/so2-pollution>.

- Ferrer-i Carbonell, A. (2005). Income and well-being: An empirical analysis of the comparison income effect. *Journal of Public Economics* 89(5-6), 997–1019.
- Forget, E. (2011). The town with no poverty: The health effects of a Canadian guaranteed annual income field experiment. *Canadian Public Policy* 37(3), 283–305.
- Fortin, N. and T. Lemieux (2015). Changes in wage inequality in Canada: An interprovincial perspective. *Canadian Journal of Economics* 48(2), 682–713.
- FRED (2018). Spot crude oil price: West Texas Intermediate (WTI). Retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/WTISPLC>.
- Gerdtham, U.-G. and M. Johannesson (2004). Absolute income, relative income, income inequality, and mortality. *Journal of Human Resources* 39(1), 228–247.
- Ifcher, J., H. Zarghamee, and C. Graham (2018). Local neighbors as positives, regional neighbors as negatives: Competing channels in the relationship between others' income, health, and happiness. *Journal of Health Economics* 57, 263–276.
- Kilian, L. and X. Zhou (2018). The propagation of regional shocks in housing markets: Evidence from oil price shocks in Canada. *Bank of Canada Staff Working Paper 2018-56*, 59 pp.
- Kronenberg, C. and J. R. Boehnke (2019). How did the 2008-11 financial crisis affect work-related common mental distress? Evidence from 393 workplaces in Great Britain. *Economics Human Biology* 33, 193–200.
- Lleras-Muney, A. (2018). Mind the gap: A review of the health gap. *Journal of Economic Literature* 56(3), 1080–1011.
- Luttmer, E. F. P. (2005). Neighbors as negatives: Relative earnings and well-being. *Quarterly Journal of Economics* 102(3), 963–1002.
- Marchand, J. (2012). Local labor market impacts of energy boom-bust-boom in western Canada. *Journal of Urban Economics* 71(1), 165–174.
- Marchand, J. (2015). The distributional impacts of an energy boom in western Canada. *Canadian Journal of Economics* 48(2), 714–735.
- Marchand, J. and J. Weber (2018). Local labor markets and natural resources: A synthesis of the literature. *Journal of Economic Surveys* 38(2), 469–490.
- Marmot, M. (1999). Multi-level approaches to understanding social determinants. In L. Berkman and I. Kawachi (Eds.), *Social epidemiology*. Oxford, UK: Oxford University Press.
- Marmot, M. (2004). *The status syndrome: How social standing affects our health and longevity*. New York: Times Books.
- Marmot, M. and R. G. Wilkinson (2005). *Social determinants of health*. Oxford: Oxford University Press.

- McEwen, B. S. (1998). Protective and damaging effects of stress mediators. *New England Journal of Medicine* 338, 171–179.
- McLinden, C. A., V. Fioletov, K. F. Boersma, N. Krotkov, C. E. Sioris, J. P. Veefkind, and K. Yang (2012). Air quality over the Canadian oil sands: A first assessment using satellite observations. *Geophysical Research Letters* 39(4), 8 pp.
- Miller, D. L. and C. Paxson (2006). Relative income, race, and mortality. *Journal of health economics* 25(5), 979–1003.
- Milligan, K. and M. Stabile (2009). Child benefits, maternal employment, and children’s health. *American Economic Review: Papers and Proceedings* 99(2), 128–132.
- Milligan, K. and M. Stabile (2011). Do child tax benefits affect the well-being of children? *American Economic Journal: Economic Policy* 3(3), 175–205.
- Poloz, S. S. (2015). Riding the commodity cycle: Resources and the Canadian economy. Speech to Calgary Economic Development, Calgary.
- Ruhm, C. J. (2000). Are recessions good for your health? *The Quarterly Journal of Economics* 115(2), 617–650.
- Ruhm, C. J. (2015). Recessions, healthy no more? *Journal of health economics* 42, 17–28.
- Schlenker, W. and W. R. Walker (2016). Airports, air pollution and contemporaneous health. *Review of Economic Studies* 83(2), 768–809.
- Schneiderman, N., G. Ironson, and S. D. Siegel (2005). Stress and health: Psychological, behavioural and biological determinants. *Annual Review of Clinical Psychology* 1, 607–628.
- Schwandt, H. (2018). Wealth shocks and health outcomes: evidence from stock market fluctuations. *American Economic Journal: Applied Economics* 10(4), 349–377.
- Smith, J. P. (1999). Healthy bodies and thick wallets: The dual relationship between health and economic status. *Journal of Economic Perspectives* 13(2), 145–166.
- Strumpf, E. C., T. J. Charters, S. Harper, and A. Nandi (2017). Did the great recession affect mortality rates in the metropolitan united states? effects on mortality by age, gender and cause of death. *Social Science & Medicine* 189, 11–16.
- Veblen, T. (1899). *The theory of the leisure class*. New York: MacMillan.
- WHO (2008). *Closing the gap in a generation: Health equity through action on the social determinants of health*. Geneva: Commission on social determinants of health - World Health Organization.
- Wilkinson, R. G. (1996). *Unhealthy societies: the actions of inequality*. London: Routledge.
- Wilkinson, R. R. (1998). Low relative income affects mortality. *British Medical Journal* 316(7144), 1611.

Appendix

Table A.1: Occupational Concentration of Oil Sector Workers

	All		Men		Women	
	Oil	Other	Oil	Other	Oil	Other
Senior management	0.54	0.52	0.53	0.70	0.57	0.34
Other management	5.04	7.48	5.12	8.64	4.66	6.30
Professionals in business and finance	4.16	2.91	2.06	2.47	13.20	3.37
Financial, secretarial and administrative	4.75	5.18	1.31	1.84	19.61	8.60
Clerical	6.64	11.94	2.33	6.54	25.19	17.47
Natural and applied sciences	15.87	7.78	15.58	11.83	17.14	3.64
Health, nurse supervisors and registered nurses	0.03	2.73	0.02	0.63	0.06	4.88
Technical, assisting and related occupations in health	0.16	3.81	0.13	1.05	0.32	6.64
Social science, government service and religion	0.99	4.53	0.61	2.39	2.60	6.72
Teachers & professors	0.26	4.73	0.27	3.04	0.20	6.46
Occupations in art, culture, recreation and sport	0.3	2.07	0.10	1.82	1.17	2.32
Wholesale, technical, insurance, real estate sales specialists	0.88	2.86	0.92	3.35	0.72	2.35
Retail Sales	0.01	5.65	0.01	3.69	0	7.65
Chefs and cooks	0.24	3.12	0.18	2.45	0.48	3.80
Occupations in protective services	0.24	1.75	0.21	2.74	0.34	0.74
Childcare and home support workers	0	1.44	0	0.20	0	2.72
Sales and service	0.73	7.47	0.48	6.05	1.80	8.92
Contractors in trades and transportation	1.31	1.08	1.57	2.02	0.21	0.12
Construction trades	1.55	2.34	1.83	4.51	0.34	0.13
Other trades	11.24	6.03	13.65	11.45	0.85	0.48
Transport and equipment operators	5.6	3.8	6.62	6.81	1.19	0.71
Trades helpers, construction, and transportation labourers	0.87	2.62	0.96	4.48	0.44	0.73
Primary industry	34.22	1.63	40.36	2.48	7.76	0.76
Machine operators and assemblers in manufacturing	4.32	5.23	5.08	7.33	1.05	3.07
Labourer in processing, manufacturing and utilities	0.07	1.29	0.06	1.50	0.11	1.09
Observations (unweighted)	70,644	2,281,785	59,744	1,122,264	10,900	1,159,521


Notes: Weighted proportions. Sample of labour force participants aged 20-60, excluding full-time students.
Source: Labour Force Survey 2006-2010, PUMFs, Statistics Canada. NOC-S 2006 codes in parentheses.

Table A.2: Probability of Hospitalization (absolute income analysis)

	Full sample		Men only	
Δp_t	0.0000	0.0000	0.0003	0.0003
(κ_1)	(0.0000)	(0.0000)	(0.0002)	(0.0002)
Oil_i	-0.0010***		-0.0005*	
(κ_2)	(0.0003)		(0.003)	
$\Delta p_t \times Oil_i$	-0.0010**	-0.0010**	-0.0022**	-0.0022**
(κ_3)	(0.0004)	(0.0004)	(0.0009)	(0.0010)
Male Dummy	-0.00031**			
	(0.00014)			
Age 25 to 34	-0.0104***		-0.0121***	
	(0.0002)		(0.0002)	
Age 35 to 44	-0.00776***		-0.0100***	
	(0.0001)		(0.0003)	
Age 45 to 54	-0.00482***		-0.0063***	
	(0.00011)		(0.0002)	
Married	-0.00068***		-0.0003*	
	(0.0001)		(0.0001)	
Single Parent	0.0004***		0.0004*	
	(0.000117)		(0.000190)	
Num. of people in HH	-0.0003***		-0.0002***	
	(0.00002)		(0.00002)	
High school degree holder	-0.0019***		-0.0018***	
	(0.00026)		(0.00026)	
University degree holder	-0.0035***		-0.0035***	
	(0.00037)		(0.00036)	
Trade or college degree holder	-0.0016***		-0.0016***	
	(0.0002)		(0.00018)	
Employed	-0.0036***		-0.0047***	
	(0.0004)		(0.00044)	
Salaried worker	-0.0015***		-0.0014***	
	(0.00026)		(0.00027)	
HH Income	0.0000***		0.0000***	
	(0.0000)		(0.0000)	
Urban residence indicator	0.0000		0.0000	
	(0.00017)		(0.00025)	
Immigrant indicator	-0.00226***		-0.00188***	
	(0.000140)		(0.000111)	
CMA/CA Fixed Effects	✓		✓	
Individual Fixed Effects		✓		✓
Observations	23,310,670	23,310,670	11,334,360	11,334,360


Notes: Statistical significance: * 10% ** 5% *** 1%. Standard errors are clustered at the CMA/CA level. All columns include controls for immigrant status, single parent, marital status, number of individuals in the economic family (log), income and income² (log), as well as CMA/CA, quarter and year fixed effects. Δp_t is scaled up by 100.

Figure A.1: Non-oil workers: Impact of changes in oil prices on the probability of hospitalization (β_1 estimates for all men and women, across local concentration of oil workers)




Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, marital status, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.

Figure A.2: Non-oil workers: Impact of changes in oil prices on the probability of hospitalization (β_1 estimates for all men and women, across local concentration of oil workers)




Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, marital status, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.

Figure A.3: Oil workers: Impact of changes in oil prices on the probability of hospitalization (β_3 estimates from inverse hyperbolic sine transformation, across local concentration of oil workers)




Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, marital status, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.

Figure A.4: Non-oil workers: Impact of changes in oil prices on the probability of hospitalization (β_1 estimates from inverse hyperbolic sine transformation, across local concentration of oil workers)




Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, marital status, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.

Figure A.5: Nursing home and hospital workers: Impact of changes in oil prices on the probability of hospitalization
 (β_3 estimates, across local concentration of nursing home and hospital workers)




Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, marital status, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.

Figure A.6: Non-nursing home and hospital workers: Impact of changes in oil prices on the probability of hospitalization
 (β_1 estimates, across local concentration of nursing home and hospital workers)




Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, marital status, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.

Figure A.7: Professional, scientific and technical services workers: Impact of changes in oil prices on the probability of hospitalization (β_3 estimates, across local concentration of professional/scs/tech workers)



Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, marital status, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.

Figure A.8: Non-professional, scientific and technical services workers: Impact of changes in oil prices on the probability of hospitalization
 (β_1 estimates, across local concentration of professional/scs/tech workers)



Note: 95% confidence intervals, standard errors clustered at the CMA/CA level. Coefficients from regressions controlling for immigrant status, single parent, marital status, number of individuals in the economic family, income and income², as well as CMA/CA, quarter and year fixed effects.