Economy Wide Spillovers From Booms: Long Distance Commuting and the Spread of Wage Effects☆

David A. Green[a], Rene Morissette[b], Ben M. Sand[c]

[a] *University of British Columbia and Research Fellow, IFS, London*
[b] *Statistics Canada*
[c] *York University*

Abstract

Since 2000, US real average wages have either stagnated or declined while Canadian average wages increased by almost 10%. We investigate the role of the Canadian resource boom in explaining this difference. We construct a model of wage setting that allows for spillover effects of a resource boom on wages in non-resource intensive locations and formulate an empirical specification based on that model. A key feature of this (and other) resource booms was the prevalence of long distance commuting - working in a resource location but residing in another community. The core idea in our model is that the expansion of the value of the commuting option during the boom allowed non-commuters to bargain higher wages. We find that wages do rise in areas with more long distance commuting. Combining these spillover effects with bargaining spillover effects in resource boom locations, we can account for 49% of the increase in the real mean wage in Canada between 2000 and 2012. We find similar effects of long distance commuting on wages in the US but the resource boom was less salient in the US and the effect on wages was one-tenth of that in Canada. Our results have implications for other papers measuring the impacts of resource booms on wages in surrounding areas. Our main finding is that long-distance commuting can integrate regions in a way that spreads the benefits and costs of a boom across the economy.

Keywords: Wages, Resource Boom, Inequality

1. Introduction

Since 2000, the Canadian and US labour markets have diverged substantially. The post-2000 period has seen stagnation or decline in wages in the US, with the real average hourly wage for prime age workers falling by 2 to 3 percent for males and being generally flat for

☆We thank Paul Beaudry, Thomas Lemieux, and participants at the Small Differences that Matter Conference, the 2017 SOLE and 2017 CEA Conferences, and at seminars at McMaster University and Dalhousie University.

Email addresses: david.green@ubc.ca (David A. Green), rene.morissette@canada.ca (Rene Morissette), bmsand@yorku.ca (Ben M. Sand)

females between the early 2000's and 2015. In contrast, in Canada, the real hourly wage rose by nearly 10% for males and 15% for females over the same period. Most strikingly, US wages declined substantially in the years just after the 2008 recession while Canadian wage series show only a short break in their upward trend at that point. Similarly, employment rates did not fall as much in Canada as the US. The term "Great" really did not apply to the 2008 recession in Canada. The two countries also followed different paths for wage inequality in this period. The log 90-10 wage differential rose by over 10% for both men and women in the US while in Canada, that differential was essentially flat, especially after 2003. Thus, the real wage growth in Canada was shared across the distribution in Canada while wages at the bottom of the US distribution, in particular, were falling.

In this paper, we seek to understand why Canada and the US have followed such different wage paths since about 2000. We argue that a substantial part of the difference can be explained by the impact of the resource boom that started in about 2003 for the Canadian economy. For anyone living in Canada during this period, such a claim would not seem surprising. The resource boom held a salient place in discussions about the economy and also in policy making in this period. But to an empirical economist who has worked with shift-share type decompositions, the claim might seem unlikely to be true. This is the case because the extractive resource sector (mining plus oil and gas) makes up only 1.5% of employment in Canada in this period. This is triple the proportion for the US and reflects some substantial growth during the period, but working with such small proportions, it is not possible to generate substantial changes in the average wage for the economy as a whole. Since the standard 'Between' component in a shift-share composition is obtained by multiplying the change in the proportion of workers in a sector by the wage premium paid in that sector, the implied component must necessarily be very small. Given results in Black et al. (2005) showing that direct spillovers to demand for other products produced locally are small, the implied effects would remain small even after taking account of derived demand effects. Our core idea, following on Beaudry et al. (2012), is that changes in the size of and wages in a salient sector can have far-reaching spill-over effects through bargaining. When a high paying sector such as extractive resources expands or starts paying higher wages, the outside option for workers in other jobs improves. In bargaining with their employers, they can point to the improved employment conditions in the resource sector and credibly threaten to quit to get a resource sector job unless their current employer increases their wages.[1] Importantly, this threat can be used by workers across the economy at the same time, resulting in a substantial multiplication of the direct wage effects of the resource boom. In essence, there is no longer a clean separation of the Between and Within components of the shift-share decomposition and so the standard between component does not provide an accurate picture of the full impact of the boom. Indeed, in Beaudry et al. (2012)'s work

[1]In an article about the industrial heartland of the US, the Globe and Mail newspaper tells the story of Cathy McClure who lost a good paying job in a forklift factory at the time of the 2008 recession. Recently, the firm returned to operation and offered her an entry level position at $15 per hour. Rather than moving to that job, she negotiated a wage increase at her current employer. Thus, the introduction of a good paying job would have benefited both whoever actually took the job and, through a bargaining effect, Cathy McClure Slater (2016).

with US cities, the spillover effect of shifts toward a higher paying sectoral composition on the average wage in a local economy is about 3 times the standard Between effect on its own.

What makes the resource boom a plausible source of these wage spillovers (apart from its salience in general discussion and policy circles) is the movement of wages in the three provinces with the highest concentration of their employment in the extractive resource industries - Alberta, Saskatchewan, and Newfoundland. The average wage in those provinces rises sharply through the mid-2000s and then grows somwehat more slowly after the recession - a pattern matching the oil price series. Fortin and Lemieux (2015) and Green and Sand (2015) both argue that regional data patterns for Canada fit well with the idea of the resource price boom driving regional differences across Canada. We replicate some of those results in the second section of this paper. Both Fortin and Lemieux (2015) and Marchand (2015) present evidence that the oil boom had impacts in other sectors of the Alberta economy. Working with a specification similar to ours, Fortin and Lemieux (2015) show that wage spillovers to other sectors from the resource sector can account for over half the difference in wage growth between oil rich Alberta and more industrial Ontario in this period. Our work complements theirs in looking beyond the borders of the extractive resource intensive provinces to ask whether the resource boom had even broader effects.

At the heart of our approach is a rich administrative data set. We use the universe of Canadian individual tax filers between 2000 and 2012. Importantly for our purposes, we have information from their main tax form (the T1), including the address from which they filed their taxes, as well as their job specific tax forms (T4's) on which is recorded the province of location of the firm for which they worked on that job. Using this, we can identify what we call long distance commuters to the extractive resource provinces: workers who file their taxes in a non resource intensive province but worked their main job in a resource intensive province. Working at the level of Statistics Canada economic regions (which roughly correspond to commuting zones), we examine the wage changes of what we call residents (workers who do not migrate or commute to other provinces to work) in local labour markets in the non-resource provinces to see whether those changes were different in locations with more commuting to resource provinces. In other words, we investigate whether increases in the value of the commuting option led to increases in wages for non-commuters. There could be wage effects on residents because of the type of bargaining effects mentioned earlier but also from supply effects, as the commuters and migrants leave the local labour force, and demand effects, as the commuters bring their earnings back to their home communities. We can eliminate the latter two channels and focus on bargaining effects by controlling for the employment rate. We present specifications with and without employment rate controls in order to see the total spillover effects and to isolate the size of bargaining effects. Those specifications indicate that about $\frac{5}{6}$ of the spillover effects we estimate occur through the bargaining channel.

Estimation of these spillover effects faces two main identification challenges. The first is that commuting may be correlated with unobserved productivity shocks to local wages. For example, if commuting to resource intensive provinces increased the most in areas that were falling into a local recession then there would be a built-in negative bias to the estimated effects of commuting on wages. We set about addressing these concerns by first constructing a model of wage setting with multiple sectors and two regions (one with a resource boom and the other without). The model follows Beaudry et al. (2012) and Beaudry et al. (2014) and, like them, we use the theory to establish the form of our estimating equations, the content of the error term (and, from it, the nature of the endogeneity problems), and candidates for instrumental variables. We construct Bartik-style instrumental variables which rely on a combination of migration links across regions in 2000 and changes in average wages in the resource provinces driven by oil sector changes. We note that over-identifying restrictions specified by the model and related to the instruments are not rejected in this data.

The second identification challenge stems from potential selection of migrants and commuters. If, for example, lower ability workers take up the commuting option because they are the ones who do not have stable employment at the time of the resource boom then we could see an increase in average wages in our local non-resource economies just because of a change in the composition of workers. We address this in some specifications in which we restrict our sample to individuals who are present in non-resource province communities throughout our data period from 2000 to 2012. That group of non-movers is very likely highly selected but since it is the same set of people throughout, any changes in their wages will not stem from composition shifts.

Our estimates indicate that there were substantial spill-overs to wages in non-resource provinces from the resource boom. In Cape Breton - an economic region on Canada's east coast that was one of the largest senders of commuters to the oil patch - we estimate that the increase in the value of the option of commuting to a resource intensive province can account for two-thirds of the 19% rise in the mean real wage for all workers that occurred in that community between 2000 and 2012. On the other hand, Toronto - which had a very low rate of commuting to resource provinces - had almost no wage spill-over effect. Taken together, we show that these spill-overs can account for about a quarter of the increase in mean wages in the Maritime provinces in our time period, with less in other provinces. We also estimate the size of bargaining related spillovers within the ER provinces themselves (which are large) and spillovers through demand by the ER sector for the products of industries in non-ER provinces (which are small). Adding all these effects together, we find that the boom can explain 49% of the increase in mean wages in Canada after 2000. This is almost certainly a lower bound on the full effect since, for example, it leaves out increases in public sector hiring that resulted from increased government revenues from the boom.

In the final section of the paper, we also provide estimates of commuting related spillover effects using US Census data. While the data is not as clean as our Canadian tax data, the estimated spillover effects are remarkably similar in the two countries. However, the growth of the ER sector relative to the rest of the economy and the proportion of workers taking up the commuting option were much smaller in the US, implying effects of the boom in resource prices in the US that are about one-tenth of what we estimate for Canada. We also construct estimates of commuting related spillover effects from the pre-recession housing boom, finding that those effects are also substantial. This suggests that the masking effects of that boom described by Charles et al. (2016) may have been larger than initially suspected.

We see these results as potentially interesting for several reasons. First, our results reveal regional interactions beyond simple migration that provide new insights into how regional economies are connected. In our model, the fact that the wage gains of permanent migration are offset by house price increases implies that permanent migration should have no bargaining effect on wages in the locations from which the migrants emigrated in equilibrium. Since long distance commuting does not have house price effects of this type and has associated frictions that are similar to finding a job in a local economy, it is predicted to have more substantial spill-over effects. We find both that permanent emigration has no effects on local wages and that increases in the value of commuting does. The fact that changes in outside options alter wages in other sectors even when holding the employment rate constant fits with recent findings pointing toward rents as important components of wages (Green, 2015). For example, in findings that firm specific effects are important components of wages in matched firm-worker data (e.g., Card et al. (2013)), our results indicate that one potential determinant of a high paying firm is whether it is located in a city with strong links to other high paying regions.

There is a growing literature on the effects of resource booms on local economies which tends to find that those booms have disproportionate effects on wages and employment relative to the size of the resource sector (e.g., Black et al. (2005),Michaels (2011), Marchand (2012), Weber (2012), Jacobsen and Parker (2016), Allcott and Keniston (2014), Feyrer et al. (2017) and Bartik et al. (2017). Some of these studies find, further, that there are positive spillover effects on production in other sectors and attribute both these effects and some of the wage effects to agglomeration externalities (e.g., Michaels (2011)). For the most part, however, explanations for the spillovers are speculation. Our model and associated results provide an alternative explanation for some of the wage effects that are found across the local resource boom economy: that they reflect bargaining spillovers in a broad set of sectors not directly related to the resource sector. Moreover, our commuting related results indicate that those spillovers extend across the whole economy. Beyond being interesting in its own right, this finding has implications for estimates in other papers since those estimates are obtained by comparing outcomes in resource boom regions with other regions. Even in Feyrer et al. (2017) and Allcott and Keniston (2014), where the authors allow for spillovers up to 100 and 400 miles from resource locations respectively, distant regions are treated as controls. Our results indicate that some of those comparison areas are being affected by the boom through long distance commuting.

Our results are also potential interesting because they help in understanding the evolution of the Canadian wage structure and the nature of the impact of the resource boom on the Canadian economy as a whole. This is important in the Canadian context because there is an ongoing debate about diversification and potential over-reliance on the resource sector in Canadian economic growth. We find that the resource boom had effects across a wide part of the economy, including geographically distant regions. That implies wage benefits for workers but also has the potential to reduce employment in those other regions. In essence, firms in the non-resource regions face increased wage costs without the associated increase in demand enjoyed by firms in the resource areas. Thus, the wage spillovers have the potential to have generated a version of Dutch disease in parts of the Canadian labour market. Our results also echo the finding in Kline (2008) that wage increases lag price increases in the oil extraction sector. We find that the spillover effects on wages in other regions are lagged still further. Thus, adjustment lags could imply costs for the rest of the economy even well after the resource boom has ended.

The claim that the Canadian wage structure was strongly affected by the resource boom is also useful when considering the Canadian data point in cross-country comparisons of wage changes. Many of the explanations for movements in US wages in recent decades have centred on the impact of technology and, to a lesser extent, trade (e.g., Acemoglu and Autor (2010), Autor and Dorn (2010)). Several authors have argued that technological change impacts should show up in a similar form in the labour market outcomes in all developed economies and that, in fact, cross-country comparisons provide a means of identifying technological change effects (e.g., Antonczyk et al. (2010)). This is not necessarily the case in models of technological adoption in which the taking up of new technologies is related to an economy's relative factor supplies. Nonetheless, it seems surprising that two economies as similar as Canada and the US would have experienced such different labour market outcomes if the main driving force in both was technological change. Understanding the extent to which another force (a resource boom) drove Canadian outcomes is helpful when trying to understand what cross-country comparisons are actually telling us. In that regard, our resource boom effects have interesting parallels with the arguments in Charles et al. (2016) that the housing boom in the US masked problems related to the decline in the manufacturing sector until the collapse of that boom with the 2008 recession. In Canada, the resource boom appeared to play a similar role but with different timing since resource prices did not suffer a persistent drop at the time of the recession. They have dropped in the last two years, potentially pointing to difficulties in the years ahead for Canada.

The paper proceeds in 6 sections not including the introduction. In the first section, we set out the broad wage patterns in the US and Canada since 2000 and show that wage movements in the resource intensive provinces are quite similar to movements in the oil price. In the second section, we set out a search and bargaining model of wage setting with multiple sectors and two main regions - one with a resource boom and one without. We use that model to derive our empirical specification. In the third section, we describe the administrative data on which our main estimates are bases. In section 4, we present the results from our empirical specification and present counterfactual exercises to assess the extent of the impact of the resource boom on the Canadian wage structure. Section 6 contains some results on long distance resource commuting for the US, and section 6 concludes.

2. Core Patterns

We begin our investigation with a comparison of movements in the wage structure in Canada and the US in the last 20 years. Our focus is on the cross-country comparison. Fortin and Lemieux (2015) provide much more detail on the evolution of the Canadian wage structure at the provincial level and Green and Sand (2015) do the same at the national level for Canada for recent decades.

The Canadian data for our exercise comes from the Labour Force Survey (LFS), which is Canada's monthly representative survey for collecting labour market data. The LFS has included data on wages since 1997 and so we make that our starting point. We have data through 2015. We restrict our attention to workers who are aged 20 to 54 and are not full time or part time students in the month in order to avoid issues related to schooling and retirement. Our goal is to get series that are as close as possible to movements in the price of labour, avoiding potential composition related movements. We use the hourly wage, which is reported directly for hourly wage earners and computed for other workers. We deflate the series to 2000 dollars using the national CPI. For the US, we use the Outgoing Rotation Group sample from the Current Population Survey (CPS) for the same years. We use the same age restrictions as for the Canadian data and, again, work with real hourly wages. Further details on the data construction are given in Appendix A.

Figure 1 contains mean log hourly wage series for Canadian and US men in the left panel and Canadian and US women in the right panel. In all the figures that follow, we normalize series to 0 in 1997 in order to avoid discussions of direct exchange rate effects on wage levels. For men, the US series shows an increase of about 8% between 1997 and 2002 but a general decline after that point, with a substantial drop after 2009. The Canadian series mimics the up and down pattern of its US counterpart until about 2004 (though with more muted movements) but moves in an almost completely opposite pattern thereafter. Notably, the 2008 recession generates a stall in the overall upward pattern rather than a sizeable drop, as in the US case. For women, the picture shares the broad feature that the two countries have similar patterns up to about 2003 or 2004 (with the Canadian series being more muted) but part company thereafter. For female wages, the US series is generally flat after 2003 but the Canadian series shows strong and nearly continual growth after that point.

In figure 2, we plot the differences between the logs of the 90th and the 10th percentiles for each gender and country. In the male figure, the US 90-10 differential shows a strong upward trend, with the increases in the years just after the 2008 recession being particularly large. In contrast, for Canada, after an initial decline and rebound, the differential is quite stable. Similarly, for females, the US series shows increases in the 90-10 differential after 2004 that are not matched in the Canadian data. Underlying these patterns are quite different movements in the tails of the distribution in the two countries. In figure 3, we plot the log of the 10th percentile. For males, the pattern is very reminiscent of the mean wage plot in figure 1. For US males, the 10th percentile declines by 10% between 2003 and 2013 while the Canadian series first increases and then is flat over the same period. In the female series, also, the Canadian pattern is one of increase while the matching US pattern is one of either decline or stagnation. The 90th percentiles for the two countries (shown in Figure 4), in contrast, look much more similar, with similar long term increases in both (albeit with different timing). But this represents quite different patterns relative to the rest of the distribution. For the US, the 90th percentile is either increasing or flat while the mean and the 10th percentiles are declining - hence the strong increase in inequality in the period. For Canada, the various parts of the distribution move up together, with increases typically beginning around 2004.

The substantial differences between the two countries after 2004 seem striking to us. The

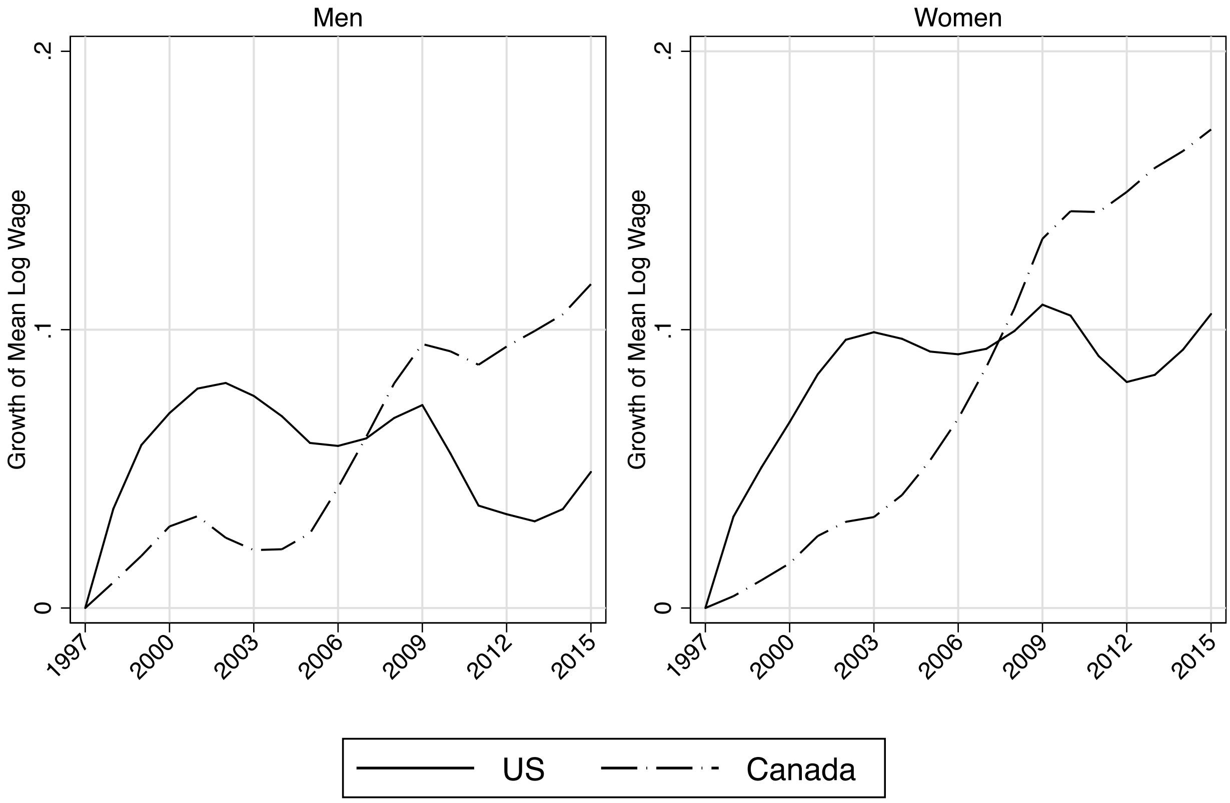

Figure 1: Real Mean Log Hourly Wages, Canada and the US, 1997-2015

two countries are similar in many regards - in the broad forms of their labour markets and labour market regulations, in their industrial structure, and in being substantial immigrant receiving nations. Most importantly, many authors have argued that technological change has been a key driver of movements in the US wage structure in recent decades and one would expect the same technological adjustments to be important in Canada. So, why have the wage patterns in the two neighbours diverged?

Both Green and Sand (2015) and Fortin and Lemieux (2015) argue that shifts in Canada's wage structure in the last 20 years have an important regional dimension. In particular, they argue that the resource boom that began in the mid-2000's had a discernible impact on wages in the three provinces with the greatest concentration in what Fortin and Lemieux (2015) call the Extractive Resource (ER) sector: Alberta, Saskatchewan, and Newfoundland.[2] To

[2]In our sample, Alberta, Saskatchewan, and Newfoundland have 8%, 4.9%, and 4.1% in the ER sector (defined as the combination of the mining and oil and gas industries), respectively. The province with the next highest ER sector employment is New Brunswick at 1.4%. Between 1997 and 2008, Alberta, Saskatchewan and Newfoundland together accounted for 96% of Canada's production of crude oil (Morissette et al., 2015).

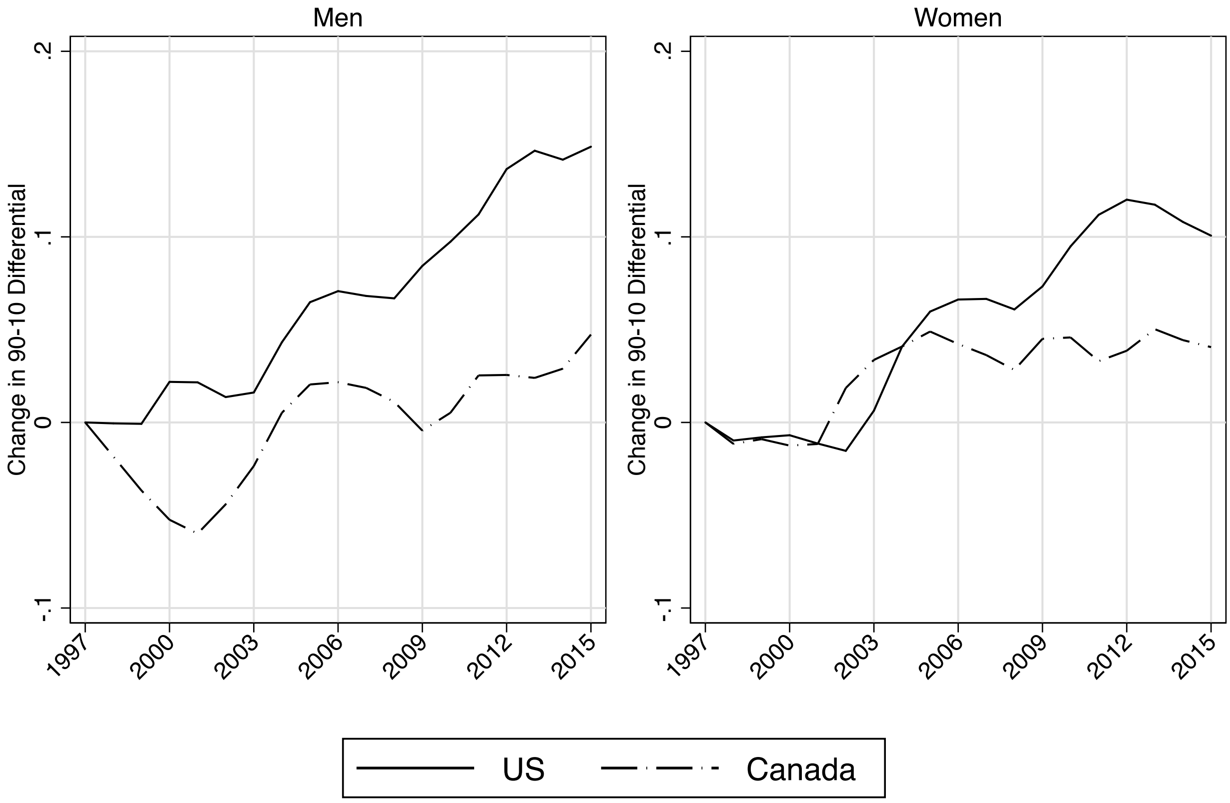

Figure 2: Log 90-10 Hourly Wage Differentials, Canada and the US, 1997-2015

illustrate this, in figure 5, we plot mean log wages for the three ER intensive provinces along with those for Ontario and for the Maritime Provinces (PEI, Nova Scotia and New Brunswick). For both men and women, the mean log wage in the ER intensive provinces increase by about 0.3 log points across our sample period, compared to wage gains of under 0.1 log points for men and about 0.1 log points for women in Ontario (the most populous province). The experience of the Maritime provinces lies between the two.

That the larger wage growth in the ER intensive provinces might be related to resource prices is supported by plots of the movement in the crude oil price (the annual average of the West Texas Intermediate price) and the mean log wage in the ER intensive provinces in figure 6. Both the wage series and the oil price series show a period of gradual growth before 2003 followed by substantial growth between about 2003 and 2009, and then a period of more gradual growth or stagnation after 2009. The timing of the inflection points in wages and oil prices are reminiscent of Kline (2008)'s findings for the US oil and gas industry that wage changes in the industry lag oil price movements by one or two years, which he attributes to workers responding to the oil price signal rather than waiting for wage signals to emerge combined with adjustment lags in labour demand. We return to this point when

Figure 3: Log 10th Percentile of the Hourly Wage Distribution, Canada and the US, 1997-2015

constructing our instruments in the next section.

The idea that the wage movements could be related to the resource price boom is also supported by the fact that the wage increases in the ER intensive provinces are much stronger among high school or less educated workers and those with some post-secondary education (including trades workers) than among those with a university education (Morissette et al., 2015). Both high school or less educated males and males with some post-secondary education experience real wage increases of over 0.3 log points in the ER intensive provinces in our sample period compared to increases of less than half that for males with a BA or more. We show figures for different education groups in Appendix A.

But, could labour market effects of the resource boom account for the differences in movements in the national wage structures in the two countries? In Canada, the ER sector makes up 1.5% of employment compared to 0.55% in the US. Alberta's 8% of workers in the ER sector compares with 2.2% in the sector in Texas.[3] Thus, the resource sector is more salient in Canada. Nonetheless, one might doubt that one sector could cause such large

[3] Wyoming is the only state with an ER employment share above Alberta's at 12.1% but Wyoming makes up only 0.18% of US employment compared to Alberta's 11% of total Canadian employment.

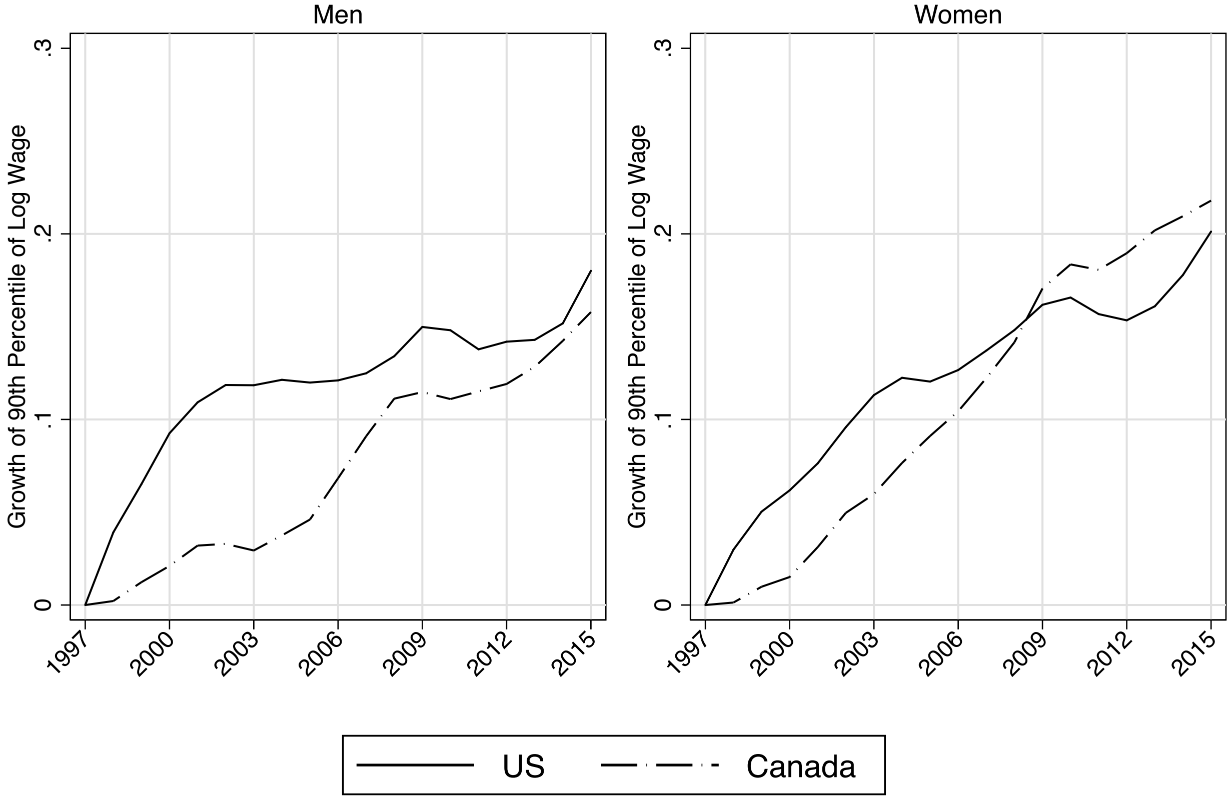

Figure 4: Log 90th Percentile of the Hourly Wage Distribution, Canada and the US, 1997-2015

changes in the average wage, even within ER provinces. After all, the ER sectors make up only 8% of Albertan employment and less in the other two ER provinces. In fact, the proportion of employment in the ER sector in Alberta increased from 0.059 in 2000 to 0.092 in 2013, and Fortin and Lemieux (2015) show that the sector paid a wage premium of 0.27 log points relative to the mean wage in the province after controlling for education, age, and gender. Combining these in a standard shift-share calculation, the increase in the size of the ER sector in Alberta would only imply a 0.9% increase in the overall mean wage for Alberta. Compared to the actual 35% increase in the mean wage in Alberta in this period, taken at face value the implication is that the increase in the resource sector during the boom had little to do with the overall increase in wages in Alberta let alone for Canada as a whole.

The obvious answer to this conclusion is that spillovers of various forms extend the impact of the resource boom beyond the small set who actually got new jobs in the sector during the boom. Papers such as Michaels (2011) and Weber (2012) hypothesize, without specific evidence, that broader effects can occur through local agglomeration externalities. Another potential channel, raised in Beaudry et al. (2012), is that increased employment and wages in the ER sector in Alberta could affect wages in other sectors through bargaining

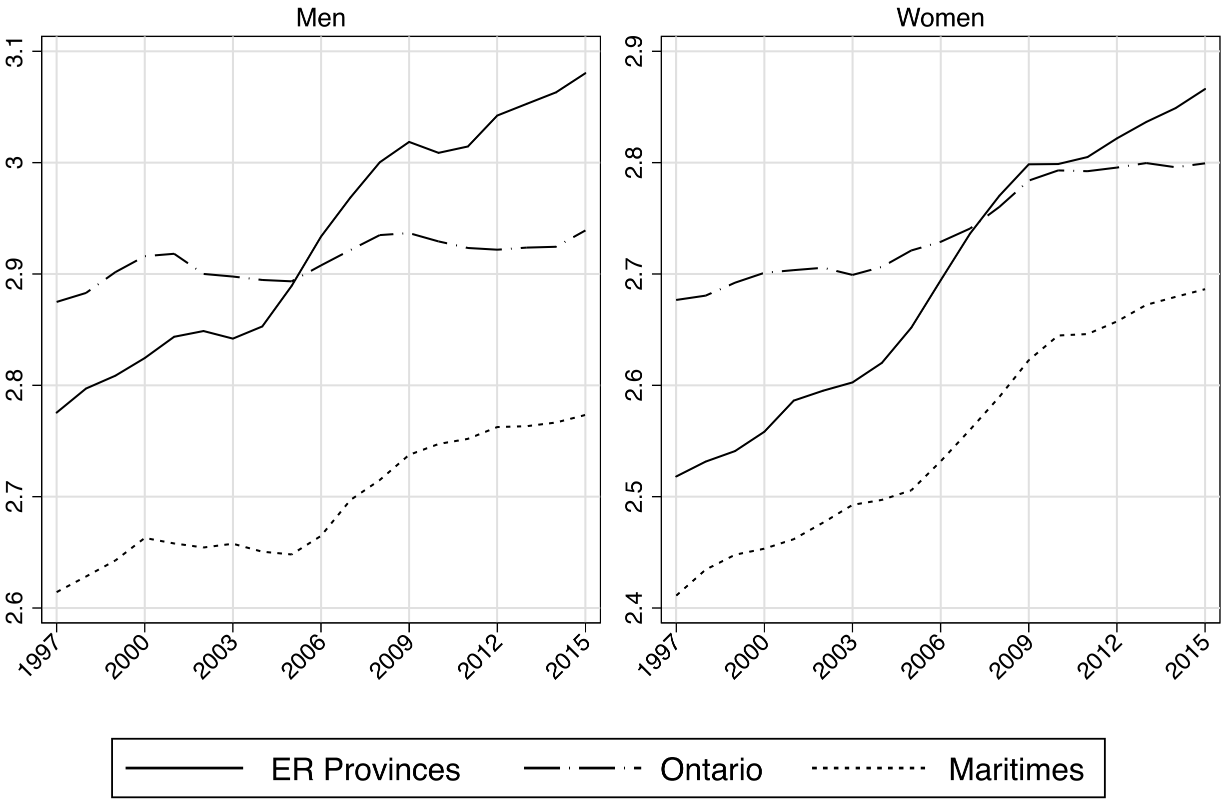

Figure 5: Mean Log Wages by Region, Canada, 1997-2015

spillovers: workers in, for example, construction in Calgary could bargain a higher wage because they had the improved outside option of going to work in the oil fields. Fortin and Lemieux (2015) examine wage changes in Alberta in this context. They implement an empirical specification that allows for bargaining spillovers and show that once those are taken into account, the ER boom can account for over half of the difference in the increase in mean wages between the ER intensive provinces and Ontario. Our argument is that the same logic may apply to places outside the ER provinces that are linked to those provinces through migration or long distance commuting and that, as a result, the ER boom could have had effects on the wage structure across Canada.

As a first check on whether geographic spillovers of this type are plausibly part of the story, we construct a counterfactual exercise in which we replace the wage increases in the ER provinces with those observed in Ontario. If the only way the ER boom affected overall wages in Canada was through changing wages and employment in the ER provinces and if the only difference between Canada and the US were due to the greater importance of the ER boom in Canada then the resulting counterfactual should look like the US. In figure 7, we re-plot the mean log wage lines for Canada and the US along with the counterfactual line for

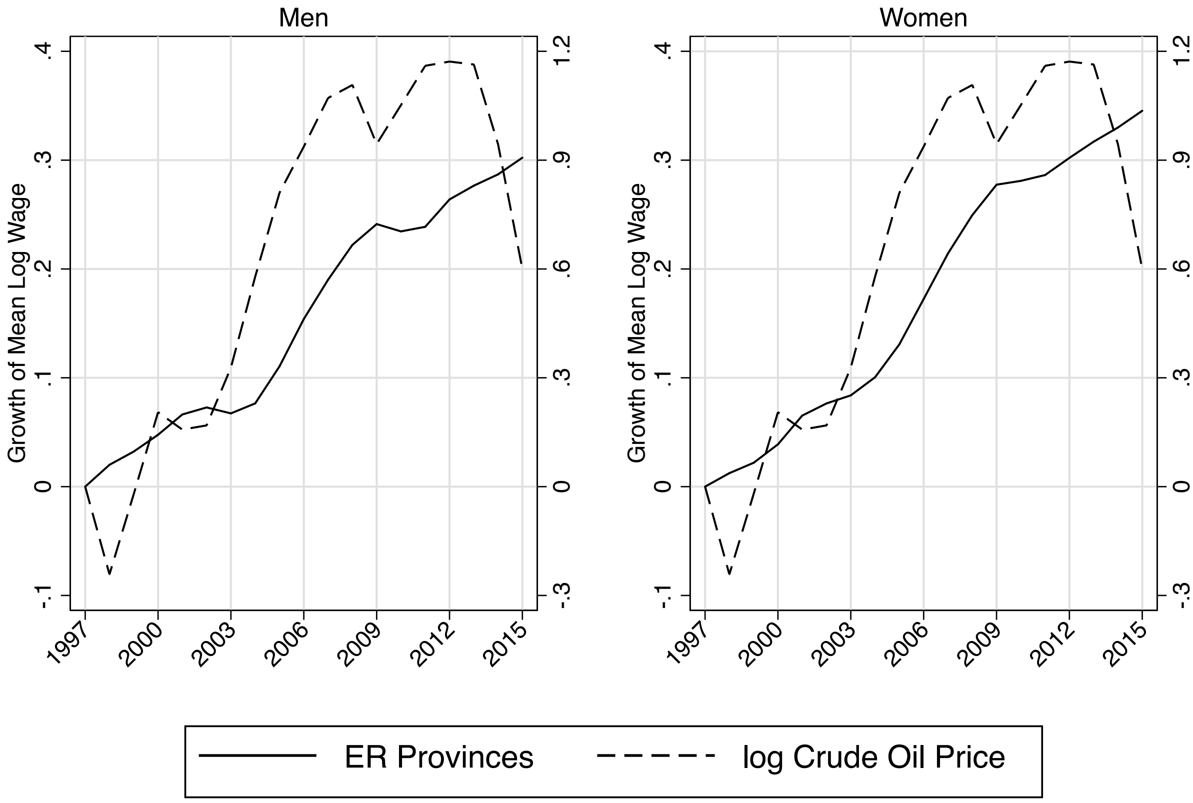

Figure 6: Crude Oil Price and Mean Log Wage in ER Provinces, 1997-2015. Left Scale: Mean Log Wage. Right Scale: Oil Price

Canada. Taking out the extra wage increases in the ER provinces mutes the national increase in Canadian wages for both men and women. For men, the increase from 1997 to 2015 is reduced from 0.11 to 0.08, which is remarkable when one notes that employment in the ER intensive provinces make up only 16% of national employment in our sample. But while taking out the wage increases in these provinces moderates the increase, it does not change the overall pattern in which mean wages surge around the time of the oil price increase and do not decline persistently after the recession as they do in the US. Put differently, wages in other parts of Canada also have a time pattern that echoes resource price changes. Given the low proportions of workers in other provinces in the ER sector (only 0.6% in Ontario, for example), this is suggestive of wage spillovers from the resource sector in the ER intensive provinces to wages in other provinces.

The decline in male wages in the US that starts just after 2000 and accelerates during and just after the 2008 recession is reminiscent of the argument in Charles et al. (2016) that a decline in manufacturing employment in the US that starts at least as early as 2000 was partially masked by demand from the construction industry during the housing boom. With

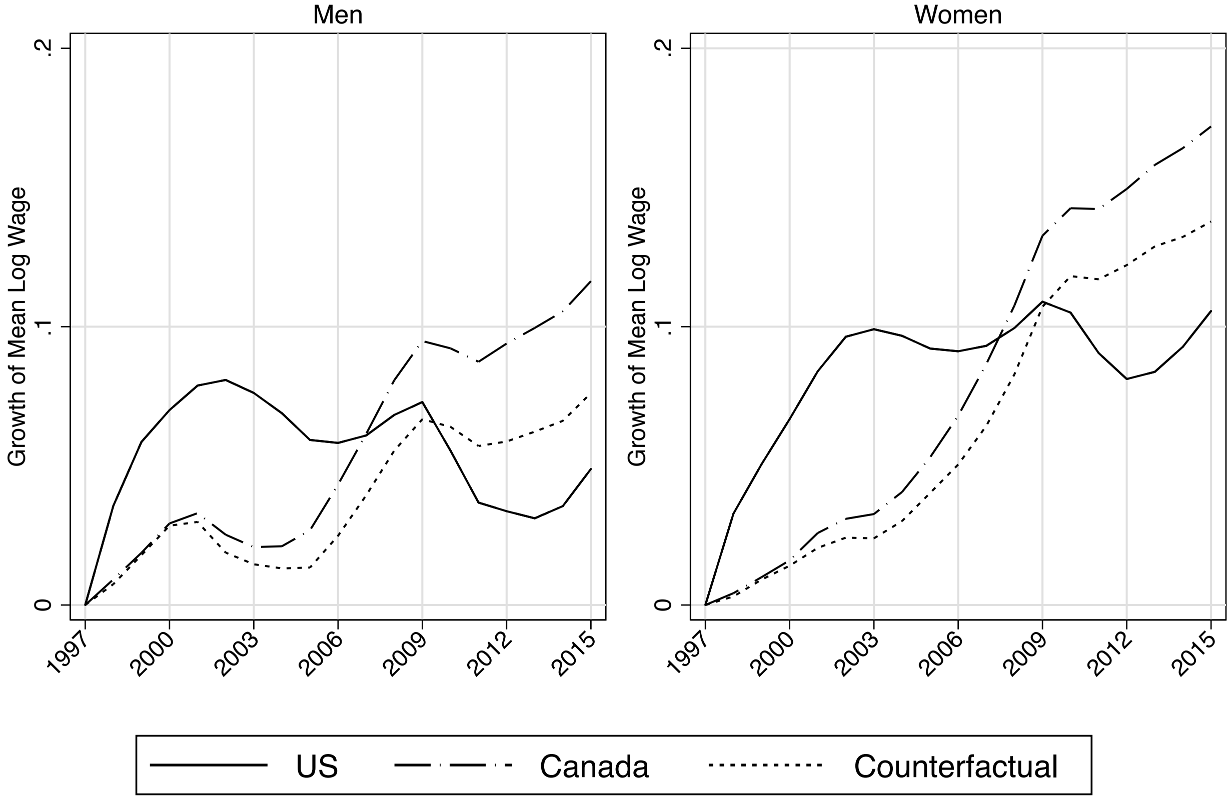

Figure 7: Mean Log Wages with No-ER Province Wage Increase Counterfactual, 1997-2015

the collapse of that boom during the recession, conditions in the male labour market also collapsed and the lack of alternative jobs in the manufacturing sector was made apparent. We believe that the resource boom may have played the role of masking the decline in manufacturing employment in Canada not only in the resource rich provinces but in other provinces as well.[4] Since resource prices only briefly declined and then recovered during and after the recession, Canadian wages and employment did not suffer the same drops after 2008 as in the US. However, we do not view it as the sole determinant of movements in the Canadian wage structure and their differences relative to the US. Fortin and Lemieux (2015) show that the minimum wage increased continually between 2005 and 2011 in the Maritime provinces and that this increase can account for much of the movement at the bottom of the wage distribution in those provinces over that period. At the other end, wage movements for those with a university degree show a smaller regional spread and less clear

[4]The decline in manufacturing employment is remarkably similar between the two countries in this period - though the precise timing is not perfectly matched. In 1999, manufacturing jobs made up just over 16% of employment in both countries. In 2015, they constituted 10.7% of US employment and 10.2% of Canadian employment.

correlation with resource price movements. To uncover the effects of the resource boom, we need an identification strategy to separate those effects from other driving forces, including movements in the minimum wage. In the next section, we set out a model of wage setting with different regions as the basis for discussing identification issues and solutions.

3. Model

In order to frame our investigation of potential labour market spill-over effects from a resource boom, we use a variant of a search and bargaining model with multiple sectors from Beaudry et al. (2012) and Tschopp (2015). In particular, we consider an environment in which the economy has two geographically separate areas. The first has oil and gas reserves and the second does not. We are interested in the wage and employment outcomes for workers in cities in the non-oil regions during and after an oil price boom. The oil boom makes it more attractive for workers from the non-oil region to migrate to work in the oil region which has two effects on wages in local labour markets. First, the actual migration of workers reduces labour supply, resulting in an increase in wages. Second, for the workers who do not migrate, the option to migrate makes their bargaining position stronger relative to their employers, also raising wages. Since many workers can present this option in bargaining without actually pursuing it, the latter, spillover type effect could be very wide ranging and large. These wage effects, to the extent they exist, could have effects on firm decisions on whether to open job vacancies and, hence, on the employment rate.

To make this more precise, consider an economy, O (Ontario), which produces a final good Y_O using output from I_O industries:

$$Y_O = \left(\sum_{i=1}^{I_O} a_i Z_i^{\chi} \right)^{1/\chi}, \qquad \text{where } \chi < 1. \tag{1}$$

where the a_i are industry specific productivity shifters and χ determines the elasticity of substitution. The price of the final good is normalized to 1, while the price of the good produced by industry i is given by p_i^O. We will assume that there are C local markets (cities) within O, with the industrial goods produced within any of them. Thus, the total amount of industrial good Z_i produced in the economy equals the sum across cities of X_{ic}, the output in industry i in city c.

Next to O is a second economy, A (Alberta), which produces a final good Y_A. Economy A differs from O in that it has oil reserves and, so, produces its output (Y_A) with $I_A = I_O + 1$ inputs, where the last input is oil. Its final good production function is otherwise the same as (1), though its final product will include oil (and could just consist of oil). An oil boom corresponds to a sudden increase in a_{I_A}, which induces an increase in p_{I_A}, the oil price. Our focus is on the direct labour market impacts in region O of that oil price increase. For that reason, we will proceed with the simplifying assumptions that the industrial goods are not traded across the two economies and that oil is not used in the production of O's output. Thus, we are not considering either the effects of oil price increases in altering the costs and, therefore, the composition of industries in O or the effects through impacts on the exchange rates between the currency of the country O + A and those of other countries . As we will discuss, though, some of our estimates will incorporate indirect effects of these types.

Much of what is interesting in the model can be learned from examining the worker and firm Bellman equations. For firms in city c (in region O) and industry, i, the value of a filled vacancy is given by

$$\rho V_{ic}^{f} = (p_i^O - w_{ic} + \epsilon_{ic}) + \delta(V_{ic}^{v} - V_{ic}^{f}), \tag{2}$$

where, ρ is the discount rate (common to workers and firms), δ is the exogenous termination rate of matches, V_{ic}^{f} is the value of a filled vacancy, V_{ic}^{v} is the value of an unfilled vacancy,and w_{ic} is the city-industry specific wage. We assume that each firm employs one worker who produces output valued at $p_i^O + \epsilon_{ic}$, with ϵ_{ic} being a city-industry specific advantage such that $\sum_c \epsilon_{ic} = 0$. Thus, the flow profits for the firm are given by $(p_i^O - w_{ic} + \epsilon_{ic})$. For the moment, we assume that all workers are homogeneous.

The expression for an unfilled vacancy is given by,

$$\rho V_{ic}^{v} = -r_i + \phi_c(V_{ic}^{f} - V_{ic}^{v}), \tag{3}$$

where r_i is the per-period cost of keeping a vacancy open and ϕ_c is the probability a firm fills a vacancy. For simplicity, and without loss of generality, we set $r_i = 0$. As in Beaudry et al. (2014), we assume that firms must pay a fixed cost, k_{ic}, to open a vacancy in industry i in city c, with that cost rising in the relative size of industry i in that city:

$$k_{ic} = \left(\frac{N_{ic}}{L_c}\right)^{\lambda} e_{ic} \tag{4}$$

where, N_{ic} is the number of vacancies (both filled and unfilled) in the i-c cell, L_c is the size of the labour force in c, $\lambda > 0$ is a parameter, and e_{ic} is a local idiosyncratic cost component to opening a vacancy. The idea behind this specification is that entrepreneurs are needed for opening a firm and exist in proportion to the population of a city. As a sector expands relative to the size of the city, it must engage less productive (higher cost) entrepreneurs. Entrepreneurs enter the sector freely until $k_{ic} = V_{ic}^{v}$. As opposed to more typical specifications where entrepreneurs are homogeneous and enter until the value of a vacancy is driven to zero, this approach permits the co-existence of different industries within each city.

Workers in the economy can be either employed or unemployed in a given period. The Bellman equation for a worker employed in industry i in city c is given by,

$$\rho U_{ic}^{e} = w_{ic} - \gamma p_{hc} + \mathfrak{A}_c + \delta(U_c^{u} - U_{ic}^{e}), \tag{5}$$

where U_c^u represents the value associated with being unemployed, p_{hc} is the price of housing, and $\mathfrak{A}_c$ is the value of local amenities in c.

The Bellman equations to this point are quite standard. The real difference in introducing sectors and regions is seen in the Bellman equation for unemployed individuals. An unemployed worker in c has a set of potential working options. The first option pertains to her initial city. With probability ψ_c, she meets a vacancy in her city, and we assume that in this case she will find jobs in proportion to the relative sizes of the industries in the city. That is, the probability that she finds a job in industry i in her home city, c, is $\psi_c \eta_{ic}$, where η_{ic} is the proportion of employment in city c that is in industry i. Conditional on not meeting a vacancy in her home city, the worker considers options in other locations. We will assume that the probability that two or more of these options are available at the same time is zero. Thus, her second option is to move to economy A, which happens with (conditional) probability μ^P_{cA}, where the P superscript refers to a permanent move. The probability, μ^P_{cA}, reflects cross-region frictions, i.e., both the difficulty of finding out about job options in the other economy and the costs of moving. We assume that once in A, the person joins the pool of unemployed workers there and searches for a job. The third option is for the person to work in A but live in c - an option that became more common during the height of the oil boom.[5] We call the probability of a worker taking that option, μ^T_{cA}, where the T stands for temporary. The fourth and fifth options are the same as options two and three except they refer to getting a job in a different city in Ontario. Calling this other city, b, we have probabilities of moving to the other city permanently and undertaking long distance commuting to it as μ^P_{cb} and μ^T_{cb}, respectively.

Given these options, we can write the value associated with being unemployed as,

$$\rho U^u_c = d - \gamma p_{hc} + \mathfrak{A}_c + \psi_c \cdot \left(\sum_j \eta_{jc} U^e_{jc} - U^u_c \right). \tag{6}$$

$$+(1-\psi_c)[\mu^P_{cA}(U^u_A - \theta_{cA} - U^u_c) + \mu^T_{cA}(U^e_{cA} - \tau_{cA} - U^u_c)$$
$$+\sum_b [(\mu^P_{cb}(U^u_b - \theta_{cb} - U^u_c) + \mu^T_{cb} \sum_j (U^e_{cb} - \tau_{cb} - U^u_c)]]$$

where, d is the flow value of unemployment, U^u_A is the value of unemployment for a person living in A, θ_{cA} is the cost of making a permanent move from c to A, U^e_{cA} is the value of a commuting job in A for a person who remains living in c, τ_{cA} is a fixed cost of taking up the option of commuting to A, and U^u_b, U^e_{cb}, θ_{cb}, and τ_{cb} are defined analogously for economies, b, in region O. Note that we are assuming that individuals can only search for new options while unemployed, which allows for a transparent empirical specification, as we will see.

It is useful to discuss each of the worker options in more detail. In the home economy, the probability that vacancies and unemployed workers in city c meet is determined by a matching function, $m(L_c - E_c, N_c - E_c)$, where L_c is the size of the labour force in c, E_c is the number of employed workers (and, thus, the number of matches), and N_c is the total number of vacancies (both filled and unfilled). Thus, the number of matches is a function of the number of unemployed workers, $(L_c - E_c)$, and the number of unfilled vacancies,

[5] In 2011, what Statistics Canada calls Interprovincial Employees - workers who file their taxes in a different province from where their main employer is located - made up 3% of the paid workforce. That percentage increased by approximately a third between 2004 and 2008, with over 60% of that increase accounted for by workers doing this type of long distance commuting to Alberta.

$(N_c - E_c)$. We assume that the matching function is constant returns to scale. Given this, Beaudry et al. (2012) show that in steady state the probability a vacancy meets a worker (ϕ_c) and the probability an unemployed worker meets a vacancy (ψ_c) can both be written as a function of the employment rate, ER_c.

Equation (6) says that the value of unemployed search depends, in part, on the composition of employment in the home economy. Using (5), it is straightforward to show that $\sum_j \eta_{jc} U^e_{jc}$ can be written as a function of the average wage in the economy:

$$\sum_j \eta_{jc} U^e_{jc} = \frac{1}{\rho+\delta} \sum_j \eta_{jc} w_{jc} - \frac{\gamma}{\rho+\delta} p_{hc} + \frac{1}{\rho+\delta} \mathfrak{A}_c + \frac{\delta}{\rho+\delta} U^u_c \tag{7}$$

Workers who move to A and join the unemployed pool there face option values that are directly analogous to those faced by searchers in c. That is, they have a probability of meeting a firm, ψ_A, that can be written as a function of the employment rate, ER_A. We again assume that workers search randomly across industries, with the probability a match is in any given industry being proportional to the size of that industry. As a result, we can write:

$$\rho U^u_A = d - \gamma p_{hA} + \mathfrak{A}_A + \psi_A \cdot \left(\sum_j \eta^P_{jA} U^e_{jA} - U^u_A \right). \tag{8}$$

with terms defined analogously to (6) apart from η^P_{jA} which we write with a P superscript to allow for the possibility that the set of industrial options faced by permanent migrants differs from that for prior inhabitants.[6] Given a form for U^e_{iA} that is analogous to (5), it is straightforward to show that this can be rewritten as a function of the wages in the various industries in A:

$$\rho U^u_A = \frac{\rho+\delta}{\rho+\delta+\psi_A} d - \gamma p_{hA} + \mathfrak{A}_A + \frac{\psi_A}{\rho+\delta+\psi_A} \cdot \left(\sum_j \eta_{jA} w_{jA} \right). \tag{9}$$

Reducing expression (9) to more fundamental terms, Beaudry et al. (2012) show that in a similar set-up the average wage term, $\sum_j \eta_{jA} w_{jA}$, can be written as a linear function of the average of industry prices, $\sum_j \eta_{jA} p^A_j$. Thus, an increase in the oil price implies a higher value to relocating to A and searching for a job there. But whether changes in wages and housing prices in Alberta actually affect the value of unemployment (and, through it, the bargained wage) in Ontario is less clear cut. As in a standard Roback model, assume that housing in each location is less than perfectly elastically supplied so that inflows of workers to Alberta will drive up housing prices there and result in house price declines in Ontario. After an oil price increase, workers will move from O to A until house price differentials offset wage differentials and an equilibrium with $U^u_c = U^u_A - \theta_{cA}$ is re-established. Once that

[6] To simplify the exposition, we have written the specification for U^u_A as if the unemployed workers in A do not have the options to either migrate or commute long-distance to O.

is the case, the term related to moving to Alberta drops out of expression (6). In essence, the option of moving to Alberta is only relevant to a worker in c if there are rents to doing so. If any wage gains are exactly offset by higher house prices then the move is a matter of indifference. However, if, in the short run, wages surge ahead of local prices in Alberta then wage effects from the oil boom in Alberta will be relevant for unemployed workers in cities in Ontario. We will proceed by including the Alberta wage effects, treating the question of whether they are relevant (and whether we are in equilibrium) as an empirical matter.

We turn, next, to the long-distance commuting option. Here, we assume that unemployed individuals who do not find a job in their home market and do not get the option to permanently migrate, get a long-distance commuting job option with probability, μ^T_{cA}. Taking that job has associated with it a fixed cost τ_{cA}. The probability of getting the chance to commute varies by city c according to the strengths of connections to A that could, in turn, be a function of the number of individuals who flew back and forth between c and A in the past. As with the first two options, the opportunity to commute to A is associated with a specific sectoral distribution of potential matches. We then have:

$$U^e_{Ac} = \sum_j \eta^T_{ja} U^e_{jAc}, \tag{10}$$

with U^e_{jAc}, the value of working a commuting job in industry j given by:

$$\rho U^e_{jAc} = w_{jAc} - \gamma p_{hc} + \mathfrak{A}_c + \delta(U^u_c - U^e_{jAc}), \tag{11}$$

where, η^T_{ja} refers to the industrial composition of work for temporary (commuter) migrants.

Working from (11), we can generate a similar expression to (7), replacing $\sum_j \eta^T_{jA} U^e_{jAc}$ in (10) with a linear expression in p_{hc}, $\mathfrak{A}_c$, U^u_c and the weighted sum of wages, $\sum_j \eta^T_{jA} w_{jAc}$. This case differs from the permanent migration case in two ways. First, workers continue to face the house price and amenity costs in c. Thus, this flow is not reduced by rising housing prices in A. Second, we assume that the commuting option is like adding a set of extra industry options for searchers in the home market. As in the home market, the searchers cannot control their chance of getting an offer from this source and, as a result, it is possible that the marginal commuter may be earning a rent relative to the cost of commuting. In contrast, the permanent move option is something under the control of the searcher (apart from the frictional costs of making the move). They will pursue that option until the marginal value of doing so is driven to zero. To the extent this distinction is true, we would expect the value of commuting to affect local wages while average wages for permanent migrants in A will not have an effect. Finally, the values for permanent moves and long distance commuting to the other cities within O are written in directly analogous ways.

3.1. Deriving the Wage Equation

Once workers and firms meet, they bargain a wage according to a Nash bargaining rule:

$$(V^f_{ic} - V^v_{ic})\kappa = (U^e_{ic} - U^u_c) \tag{12}$$

where, κ is a parameter capturing the relative bargaining strength of workers versus firms. From (2) and (3), we can write:

$$(V^f_{ic} - V^v_{ic}) = \frac{p_i - w_{ic} + \epsilon_{ic}}{\rho + \delta + \phi_c} \tag{13}$$

And, from (5), we get:

$$(U^e_{ic} - U^u_c) = \frac{1}{\rho + \delta}[w_{ic} - \gamma p_{hc} + \mathfrak{A}_c] - \frac{\rho}{\rho + \delta} U^u_c \tag{14}$$

We can solve these for the bargained wage:

$$w_{ic} = \frac{(\rho + \delta)\kappa}{(\rho + \delta + \phi_c) + (\rho + \delta)\kappa}(p_i + \epsilon_{ic}) + \frac{(\rho + \delta + \phi_c)}{(\rho + \delta + \phi_c) + (\rho + \delta)\kappa}[\mathfrak{A}_c - \gamma p_{hc}] + \frac{(\rho + \delta + \phi_c)}{(\rho + \delta + \phi_c) + (\rho + \delta)\kappa} U^u_c \tag{15}$$

Then, substituting expressions for U^e_{jc}, U^u_A, U^e_{cA}, U^u_b, U^e_{cb} into (6), we get:

$$\begin{aligned} w_{ic} = {} & \alpha_0(p_i + \epsilon_{ic}) + \alpha_1 d + \alpha_2(\mathfrak{A}_c - \gamma p_{hc}) + \alpha_3(\mathfrak{A}_A - \gamma p_{hA}) + \alpha_4 \psi_c \sum_j \eta_{jc} w_{jc} \\ & + \alpha_5(1 - \psi_c)\mu^P_{cA}\psi_A \sum_j \eta^P_{jA} w_{jA} + \alpha_6(1 - \psi_c)\mu^T_{cA} \sum_j \eta^T_{jA} w_{jAc} \\ & + \alpha_7(1 - \psi_c) \sum_b \mu^P_{cb}\psi_b \sum_j \eta_{jb^P} w_{jb} + \alpha_8(1 - \psi_c) \sum_b \mu^T_{cb} \sum_j \eta^T_{jb} w_{jbc} \\ & + \alpha_9(1 - \psi_c) \sum_b \mu^P_{cb}(\mathfrak{A}_b - \gamma p_{hb}) \end{aligned} \tag{16}$$

The α parameters can be written as explicit functions of structural parameters such as δ, ρ, and ϕ_c. We do not attempt to back out estimates of the underlying structural parameters because our interest is actually in their net effects as reflected in the α's.

Equation (16) captures the idea that in a bargaining environment, the wages for a worker will depend not just on her productivity in that sector (captured by $(p_i + \epsilon_{ic})$ in this case) but also on her outside options. In our environment, those options are reflected in the average wages the individual would expect to earn from returning to unemployment and then getting another job in the local economy or getting a job in another location through either a permanent move or long distance commuting. The relative housing and amenity prices in the home location and potential permanent move locations also have an effect. As described earlier, this expression is derived assuming that we are not in a migration equilibrium. If we were then the average wages associated with a permanent move and the house prices and amenities associated with the other locations would not affect wages in c.

The fact that w_{ic} is expressed as a function of the average wage across sectors in c implies that this wage equation embodies a standard reflection problem. Partly in response to this, Beaudry et al. (2012) re-write their analogous wage equation, replacing local wages with the wage premium for industry i relative to an arbitrary base industry 0 at the national level.

That is, they show that one can derive a new wage expression in which $\sum_j \eta_{jc} w_{jc}$ is replaced with $R_c = \sum_j \eta_{jc}\nu_j$, where $\nu_j = w_j - w_0$ are the national level wage premia. We will call R_c the wage rent in city c since it shows the average wage premium across industries in c. Since we are working with homogeneous workers, these premia are interpreted as rents. We also replace the average wages in the other locations with average rent variables. For example, the average wage for permanent migrants to A, $\sum_j \eta^P_{jA} w_{jA}$, is replaced with $R^P_A = \sum_j \eta^P_{jA}\nu_j$. The explicit steps in moving from equation (16) to one in which average wages are replaced by the rent variables in a manner consistent with the model are described in Appendix B.

According to (16) wages are also a function of the employment rate in the local economy through its determination of ψ_c (which appears explicitly in (16)) and ϕ_c (which appears as parts of the α's). We want to make that relationship explicit for reasons that will become apparent in a moment. In addition, note that the average wage in A is multiplied by factors that reflect the probability of finding a job in c (ψ_c), the probability of getting the option to move (μ^P_{cA}), and the probability of finding a job in A (ψ_A). We assume that workers cannot credibly communicate their personal cost of moving when bargaining with their employer. In that case, rather than basing wage bargains on values for the cost of moving and the likelihood of finding a job in Alberta, we assume that it is based on a common estimate of the likelihood that the worker would end up with a job in Alberta if she were to leave the current match. We take that common estimate to be the proportion of workers in c who moved to A and got a job there in a recent period. That is, rather than focusing on the rent variable, R^P_A, on its own, we will work with $X^P_{Ac} = q^P_{Ac} R^P_A$, where q^P_{Ac} is the proportion of workers from c who moved permanently to A in the previous period. We take q^P_{Ac} to be used by the bargaining agents as a guess of $(1-\psi_c)\mu^P_{cA}\psi_A$, the probability an unemployed worker does not get a job in c, moves permanently to A, and gets a job in A. Similarly, we will represent the option of long-distance commuting to A with $X^T_{Ac} = q^T_{Ac} R^T_A$.

In order to make the contribution of the local employment rate explicit, we log linearize (16) around a point where all the locations in O have the same industrial composition and pay the same wage. We then difference it with respect to time in order to eliminate fixed city x industry characteristics, where different time periods correspond to different steady states. The resulting wage specification is:

$$\Delta \ln w_{ic} = \beta_{0i} + \beta_1 \Delta R_{ct} + \beta_2 \Delta X^P_{Act} + \beta_3 \Delta X^T_{Act} + \beta_4 \Delta X^P_{Bct} + \beta_5 \Delta X^T_{Bct} + \beta_5 \Delta ER_{ct} + \xi_{ict} \quad (17)$$

where $X^P_{Bct} = q^P_{Bct} R^P_{Bct}$. R^P_{Bct} is the wage rent for workers from location c who migrated to non-resource provinces. The weights q^P_{Bct} measure the fraction of workers from location c who moved to a non-resource province other than their province of residence. Likewise, $X^T_{Bct} = q^T_{Bct} R^T_{Bct}$, where q^T_{Bct} equals the fraction of workers from location c who do long-distance commuting in a non-resource province other than their province of residence. We use this instead of option values for every possible city in non-resource provinces in order to reduce our computational burden.

One important aspect of specification (17) is that the intercept has an i subscript, which reflects the fact that it contains the industry specific price terms among others. In estimation, we capture this by including a complete set of industry dummies, implying that our estimated effects are within industry over-time estimates. Intuitively, this means that we identify the local wage rent effect (β_1) by comparing the wage changes for workers in the same industry in two different cities that are experiencing different changes in their industrial composition. For example, we are identifying the bargaining effect by comparing wage changes for construction workers in Hamilton, with the loss of its high wage rent steel sector, to workers in, say, Moncton, without the loss of such a sector. The core idea is that the construction workers in Hamilton were able to bargain higher wages than those in Moncton because their outside option included the possibility of getting high wage steel jobs but they lost that advantage with the decline in the steel sector.

With the wage specification, we are in a position to discuss the potential impacts of the resource boom on communities other than those directly involved in the boom. In the model, the boom shows up as an increase in the price of oil that we treat as specifically relevant as an output in A. That increase, in turn, leads to an increase in wages in the oil producing sector that then generates increases in bargained wages in other sectors in A. Fortin and Lemieux (2015) and Marchand (2015) examine the wage spill-over effects of the oil boom in Alberta, showing that the impacts on wages in the non-oil sectors were large. Similarly, Feyrer et al. (2017) find substantial spillover effects from new fracking locations in the US on wages up to 100 miles away from those locations. Our work complements these other papers in investigating further spill-over effects to other parts of the country.

The first effect of the increase of wages in A will be to induce some workers from O to move to A and others to engage in long-distance commuting to jobs in A. The result is a reduction in the supply of workers in the communities in O, inducing an increase in wages there. In the model, this immediate impact arises because the employment rate, ER_c, rises and tighter labour markets benefit workers in wage bargaining. But without any further changes, this will induce a reduction in job vacancy creation that will eventually lead back to an equilibrium with the same wage and employment rate as before the boom. This first, labour-supply channel (whether in the short or long run) is not what is being captured in the β_2 and β_3 coefficients in (17). Because our specification includes a control for the change in the employment rate, effects that necessarily involve an accompanying change in employment (which is the case with a supply shift unless the demand curve is perfectly inelastic) are not what is being captured by our wage rent variables. We will, however, present some results from reduced form specifications in which we do not control for changes in the employment rate in order to capture the total impact of the wage changes in A, including the labour supply effects. That total impact would also include increases in demand for locally produced goods stemming from the income brought home by the commuters - another effect that is not allowed when we control for the employment rate.

Given our control for the employment rate, we interpret the β_2 and β_3 as capturing wage bargaining effects: workers in c are able to bargain higher wages because the outside option of working in Alberta has become more attractive. From our derivation of the specification based on national level wage premia, above, what we are actually estimating when we use those national level premia is the complete impact of the wage changes in Alberta taking into account all feedbacks within the c economy. As wages in construction in Alberta increase because wages in the oil sector have improved, those increased construction wages can then be used as an improved outside option for workers in other sectors, which in turn imply improved outside options for construction and oil workers, etc.. Thus, for the Rent variable for the permanent move to Alberta option, the coefficient is $\frac{\gamma_{A1}}{(1-\gamma_{A2})}[\alpha_5 + \frac{\alpha_5\alpha_4^*}{1-\alpha_4^*}]$. The first part of this expression, $(\frac{\gamma_{A1}}{(1-\gamma_{A2}})$, captures the fact that we are working with the total effect of the oil change on wages in Alberta, taking into account all feedback effects. This is then multiplied by α_5, which is the direct, first round effect of the change in the wage in A, plus $\frac{\alpha_5\alpha_4^*}{1-\alpha_4^*}$, which captures the feedback loops within the c economy. These arguments hold in the same form for the other migration/commuting options.

The effect of the A wage change on wages in c happens through two channels. The first is the change in the Rent in A, holding constant the propensity for people to move or commute to A (the q's in our X_{cA} terms). The second is due to changes in the propensity for people to move or commute to A. The A options could become more salient in bargaining in c for either reason. Finally, changes in the price of oil and/or the exchange rate could have effects by altering the industrial composition of c. For example, if both oil price and exchange rate changes are particularly bad for manufacturing then the shift in industrial composition in O away from the high paying manufacturing sector will imply declines in bargained wages in all sectors in c. Because we include ΔR_c, the change in the average wage rent in the local market, as a regressor, this last channel is not part of our estimated A wage effects. However, we will investigate a reduced form specification in which we do not include R_c to get an estimate of the net effect of the oil boom, including this channel.

Finally, it is important that the wage premia that we work with correspond to rents i.e., wage differences across industries that do not correspond to productivity differences or compensating differentials. If, instead, the wage premia corresponded to compensating differentials for elements of the work in different sectors then workers in other sectors could not use them to bargain a better wage. The higher wage in, say, the asbestos industry would just compensate for an expected loss in health and so there would be no real threat in telling your employer that you will quit to take a higher wage in the asbestos industry if she doesn't raise your wage. As pointed out in Green(2015), whether the industry premia really are rents can be tested empirically in our context: if they are not rents then the average premium should not affect wage setting within sectors.

3.2. Endogeneity and Identification

Another key aspect of the specification - and a benefit of deriving it from theory - is that we know what is in the error term. In particular, the error term has the form, $\xi_{ict} = a_1\Delta\epsilon_{ict} + a_2\Delta\sum_j \eta_{jc}\epsilon_{jct} + a_3\Delta\sum_j \eta_{jA}\epsilon_{jAt} + a_4\Delta\sum_j \eta_{jBt}\epsilon_{jBt}$. That is, it is a function of changes in the productivity shock for the specific industry-city cell and of changes in location specific averages of sector productivity shocks. The immediate implication of the content of the error term is that it implies endogeneity problems. In particular, as discussed in Beaudry et al. (2012), the ΔR_{ct} rent variable is a function of changes in the industrial composition of the local economy. That will clearly be related to the changes in the local productivity shocks and their averages that are in the error term. If we are willing to argue that the productivity shocks in Alberta are independent of those in c then the rent in A should not be correlated with the error in the wage equation for location c. However, we would expect that the propensity to move to A is correlated with local productivity shocks. In particular, we would expect larger movements to A, through either commuting or migration, if there have been negative productivity shocks in the local economy. Thus, ΔX^P_{cA}, ΔX^T_{cA}, and the comparable variables for the B locations are likely to be correlated with the error term. Finally, the employment rate in c variable, ΔER_c, is also likely to be correlated with changes in local productivity shocks and their average.

We adopt an instrumental variables approach to addressing these endogeneity issues that is similar in nature to that used in Beaudry et al. (2012). In particular, it is straightforward to show that one can decompose ΔR_{ct} as:

$$\Delta R_{ct} = \sum_j \Delta\eta_{jct}\nu_{jt-1} + \sum_j \eta_{jct}\Delta\nu_t \tag{18}$$

where, $\Delta y_t = y_t - y_{t-1}$. This decomposition says that the change in the average rent variable can be broken down into changes that arise because of changes in the industrial composition, holding the industrial wage premia constant (the first component on the right hand side) and changes that arise because of changes in the premia, holding the local industrial composition constant (the second term on the right hand side). Beaudry et al. (2012) argue that each of these components can be seen as the basis for a valid instrument. In particular, they construct a Bartik-style instrument corresponding to the first component as follows. First, construct predicted employment in each industry in economy c using the local employment level in the industry in period $t-1$ together with the national growth rate in employment in the industry, i.e., $\hat{E}_{ict} = E_{ict-1} \cdot (1 + g_{it})$, where E_{ict} is the employment level in industry i in city c at time t, and g_{it} is the growth rate in industry i employment at the national level. Using these predicted employment levels, we can construct predicted shares of employment in each industry in c in period t, $\hat{\eta}_{ict}$, and then construct an instrument as:

$$IV1_{ct} = \sum_j (\hat{\eta}_{jct} - \eta_{jct-1})\nu_{t-1} \tag{19}$$

Beaudry et al. (2012) provide a detailed discussion of the conditions under which this is a valid instrument. Here, we will just provide the intuition and direct interested readers to the proofs in Beaudry et al. (2012). First, recall that given the inclusion of industry dummies in our specification, the identifying variation we are using is across cities - within industry. That means that what we are concerned with is the cross-city correlation between $IV1_{ct}$ and ξ_{ict}. The cross-city variation in $IV1_{ct}$ comes from cross-city differences in the η_{ict-1}'s.[7] Thus, for $IV1$ to be uncorrelated with the error term, we need that cross-city differences in start-of-period industrial composition are uncorrelated with cross-city differences in the growth in city-average productivity. In other words, industrial composition at the start of a period should not be a predictor for general productivity growth in a city. This is a standard Bartik style instrument argument: start of period composition is uncorrelated with local changes in the error term. In our case, the assumption may not seem entirely credible - at least ex post, it seems possible to identify types of industrial mixes that predict general

[7]To see this, note that $\hat{\eta}_{ict}$ gets its cross-city variation from η_{ict}.

growth in an area. However, Beaudry et al. (2012) argue that the second instrument, which is constructed based on the second component in (18), provides a strong over-identification test.

The second instrument is written as:

$$IV2_{ct} = \sum_{j} \hat{\eta}_{jct} \Delta \nu_t \tag{20}$$

The cross-city variation for this instrument is also based on differences in η_{it-1} and, so, its validity is also dependent on cross-city differences in industrial composition at the start of the period being uncorrelated with changes in city-level productivity. According to the model, the two instruments should give the same estimated effects. In essence, when you bargain with your employer, it does not matter to her if your outside option has declined in value because the manufacturing sector has declined in size or because its size has not changed but the wage premium it pays has declined. On the other hand, if the core identifying assumption is incorrect then there is a correlation between the η_{ict-1}'s and the error term, and the two instruments weight that non-zero correlation very differently. Thus, if the requirement that industrial composition differences are uncorrelated with changes in city level productivity fails then we would expect the two instruments to give different answers. In $IV1$ the potentially offending correlation is weighted by national industrial growth rates and the start of period industrial wage premia, and in $IV2$ it is weighted by the growth in the industrial premia. In the US data used in Beaudry et al. (2012), the two instruments have a correlation of only 0.18, and they show that one cannot reject the over-identifying restriction at any reasonable level of significance.

We will employ these instruments for ΔR_{ct}. We construct similar instruments for each of our other rent-related variables, the X_{ct}'s. However, these instruments emphasize variation coming just from changes in average rents in A and B. As we have seen, the variation in the X_{ct} variables also comes from changes in the proportion of individuals taking up these geographic alternative options. Indeed, more of the variation in the X_{ct}'s come from changes in those proportions (what we have called the q's) than from movements in the average rents. This potentially fits with Kline (2008)'s arguments that workers respond more to the probability of getting a job as signalled by oil price movements than to wage changes. Working from that logic, we also construct an instrument that emphasizes the impact of the oil boom and which we expect to affect both the average rent and the propensity to migrate (through, for example, affecting the employment rate). In particular, we construct an instrument as the proportion of the population who moved to Alberta in 2000 times the percentage change in the oil price between t and t-1. This instrument gets across city variation for any given time period from differences in original connections to Alberta and gets over-time variation from changes in the oil price. We construct a similar instrument for the long distance commuters as the proportion of individuals undertaking long distance commuting (i.e., filing their taxes in c but being paid on the main job by a firm located in Alberta) in 2000 times the percentage oil price changes.

Finally, to address the endogeneity of ΔER_{ct}, we use a classic Bartik instrument. That is, we predict the overall employment growth in a city using the start of period industrial composition combined with national level employment growth rates. In particular, $IV3_{ct} = \sum_j \eta_{jct-1} g_{jt}$. Thus, to identify our main wage effects, we are effectively comparing two cities whose start of period composition combined with national level industrial growth imply the same expected overall growth. Then, the wage rent instruments pick up the extent to which that same overall predicted growth reflects differences in changes in the predicted industrial composition, where the composition we care about is in terms of industries ranked by their wage premia. The effects for wages in Alberta are identified by differences in start of period propensities to migrate or commute to Alberta across different cities.

4. Data

The data underlying our main regressions comes from a combination of administrative datasets. The first is the T1 Personal Master File (T1PMF), which contains the captured data from the main tax individual tax form (the T1) for all individuals in Canada.[8] From that file, we get the individual's age, gender, and postal code of residence. We work with Economic Regions as our geographic units. Economic Regions are collections of Census Districts that are at approximately the level of major cities and substantial rural areas. For example, Ontario has 11 Economic Regions with the greater Toronto, Hamilton and the Niagara Peninsula, and greater Ottawa areas being three of them. We focus our attention on the non-Extractive Resource provinces (i.e., the seven provinces leaving out Alberta, Saskatchewan, and Newfoundland), within which there are 55 Economic Regions. We work with the years 2000 to 2012 based partly on data availability and partly to focus on the years around the resource boom. We restrict our attention to tax filers who are age 22 to 64 in those years in order to minimize effects of changes in school attendance and/or retirement.

We are interested in whether people residing in the 7 non-ER provinces either worked as long-distance commuters in the ER provinces or moved to those provinces. We capture the latter by comparing the residence on the T1 in year t for a person to her residence in year t+1 and define her as a permanent mover to an ER province if the place of residence in t+1 is in an ER province. To determine long distance commuting status, we use the T4 files attached to an individual's tax filing. T4's are individual forms filed in relation to each employee by the employer. From the T4's we can determine the province of location of the employer and the worker's earnings associated with the job in the calendar year. The location on the T4 is only at the provincial level and so we define long distance commuters from Economic Region c as individuals whose residence on their T1 is in region c but whose main job (as determined by their T4 with the largest associated earnings) is at a firm issuing a T4 from an ER province. Importantly, T4's are issued at the enterprise level and, therefore, the location associated with them refers to the place of work not the location of the firm's head office. Because T4's are only issued to employees, we focus on paid workers not the self-employed, and our measures of earnings will include only paid earnings. Our main outcome variable of interest is annual earnings on the main job - the job with the highest T4 earnings in a year.

[8]The T1PMF does not include the 5% of taxfilers who file late but otherwise contains the universe of tax filers.

This has the downside of not being the price of labour (i.e., the wage) that is discussed in our model as it will include variation stemming from differences in hours worked.

We are also interested in knowing the industry of employment on the main job since our Rent variables are based on industrial composition in a location and industrial wage premia. To get industry, we link the T4's to the Longitudinal Employment Analysis Program (LEAP) file through a business number. The LEAP includes information on firms gathered from firm tax records. Finally, we also link to the Immigration DataBase (IMDB), which contains information on all immigrants arriving to Canada after 1980 to establish immigrant status (or, more precisely, whether the tax filer is an immigrant who arrived after 1980).

We are interested in wage outcomes at the local labour market (i.e., Economic Region) by industry level and, so, aggregate to that level and work in 4 year differences: 2000 - 2004; 2004-2008; and 2008-2012. Our industries are at the 3-digit NAICS level, yielding 102 industries in total. Given our 55 economic regions and our 3 four-year differences, this implies at most 16830 location by industry by year observations. We drop cells with fewer than 20 observations to reduce the impact of small cells. We have not weighted the observations in our estimation by cell size because our unit of interest is the location by industry cell and we do not want to introduce added variation that stems from a potentially endogenous variable (cell size). However, we do use robust standard errors that will, in principle, adjust for any size related heteroskedasticity in the standard errors after the estimation.

We use two estimation approaches. To understand the difference between the two, it is useful to define five types of workers: 1) Residents - individuals who reside in location c in year t and who obtain all their income from the province in which c is located[9]; 2) Commuters to ER provinces - individuals who reside in location c in year t but whose main job is in an ER province; 3) Commuters to non-ER provinces - individuals who reside in location c in year t but whose main job is in a different non-ER province; 4) Recent Migrants to ER provinces - individuals residing in an ER province in year t who resided in location c in year t-1; and 5) Recent Migrants to non-ER provinces - individuals residing in a different ER province in year t who resided in location c in year t-1. Individuals in groups 1, 2, and 3 have earnings by definition but some of the group 4 and 5 members (less than 10%) do not show earnings or employment in their new location.

We have two types of earnings variables to construct. The first is the industry rents at the national level. To obtain those, in each year (2000, 2004, 2008, and 2012) we run a regression of individual log annual earnings for all workers in all provinces on a complete set of immigrant by gender interactions, a set of quadratics in age that allow for different age profiles for each immigrant by gender group, and a complete set of industry dummies. We want to interpret the coefficients on the latter dummies as rents but are concerned that they still capture differences in productive characteristics such as education (which is not available in our administrative data). To address that issue, we also tried specifications in which we included dummies representing cells in 7 by 7 matrices defined by the 5 earnings quintiles, self-employment status, and not-employed status in year t-1 on one dimension and the same

[9] In principle, this means that some of their income could be earned in a different economic region in the same province.

7 categories in year t-2 on the other dimension. The earnings quintiles in these matrices are gender*age*province-specific, i.e. are defined for 180 groups defined by province of residence, gender, and 9 age groups (20-24, 25-29, 30-34,... 60-64). This meant that current year industry effects were calculated from differences within these groups among people who had the same broad earnings/employment/self-employment histories over the two preceding years. This is an attempt to put individuals in rough productivity cells and use only the industry variation within those cells. We present estimates where the industry differentials are estimated both with and without the inclusion of those matrix cell indicators.

Once we have the industry premia (the ν_j's), we construct a Rent variable corresponding to each of our 5 worker types in each year and location as:

$$R^k_{ct} = \sum_{i=1}^{I} \eta^k_{ict} \nu_{it} \tag{21}$$

where k indexes the 5 types, c indexes the location in the non-ER province, i indexes industry, and t indexes time. As an example, R^4_{ct}, the rent variable for permanent migrants to ER provinces, will be a weighted average of industry rents where the weights are the industrial composition of recent arrivals to an ER province from location c. In contrast, R^2_{ct}, the rent for commuters to ER provinces from c will use the industrial composition of employment in the ER provinces of people who live in c but have their main job in an ER province.

To get our dependent variable, we run individual-level regressions for residents of the non-ER provinces. In particular, we run the same regressions as described above, with the log annual earnings as the dependent variable and individual controls on the right hand side, but replace the industry dummies with a complete set of industry by location dummy variables, with associated coefficients, θ_{ic}, for the industry i, location c dummy. We keep the θ_{ic}'s corresponding to cells with at least 20 observations. Thus, our dependent variable is the mean log wage for an i,c cell purged of individual characteristic effects. As before, we do this with and without the 7 x 7 matrices representing earnings groups for the preceding two years.

In our second main estimation approach, we take advantage of the fact that we have panel data. In particular, we re-run our exercise for getting the θ_{ic}'s but keep only resident type workers who are present in non-ER provinces across all our sample years.[10] The idea is that we will look for effects on wages for a constant group of workers who never commuted or migrated to oil-producing provinces during the 2000-2012 period. This addresses specific selectivity issues, i.e. the possibility that observed resident average wages might change because the best workers have moved or started commuting to oil-producing provinces. The group who are resident in all years is, of course, a select group in their own right. But if we assume an efficiency units form for wages, in which a wage can be written as a time varying price per effiency unit times a person-specific (and constant) number of efficiency units then wage movements for a consistent set of workers (however special they may be) will reflect movements in the price of labour, which is what we care about. Because we are

[10]Our panel data sample consists of individuals who are aged 22 to 42 in 2000 (34 to 54 in 2012).

using earnings rather than wages, our estimates will include labour supply responses within the local economy by the set of consistent residents. In addition, this approach does not control for changes in selection across sectors within a location and so our measure also reflects responses to the ER price shocks of this type.

4.1. Variable variation

It is helpful for understanding the source of our identifying variation to see some of the differences in industrial composition across our different worker types and regions. In Table 1, we present the percentage of workers in 2-digit industries for our five workers types in 2000 as well as for workers who resided in ER provinces in 2000. In contrast to residents and commuters (who resided in non-ER provinces in 2000), recent migrants resided in non-ER provinces in 1999. Table 2 contains the same industrial distributions for 2012. A few key points stand out from the tables. First, for non-ER province residents the percentage in manufacturing declines from 18.5 in 2000 to 2012 in 12.0 in 2012. This pattern fits with national trends in manufacturing employment and with Charles et al. (2016)'s argument for the US that there is an underlying secular trend away from manufacturing in the US as well. Offsetting that movement are increases in construction, health, retail trade, and accommodation and food services, among others. Second, commuters to ER provinces are much more concentrated in construction, mining and oil extraction, and agriculture and forestry than their resident counterparts, and that concentration intensifies over time. For example, the percentage in mining and oil extraction approximately doubles from 8.5% in 2000 to 16.3% in 2012. Interestingly, the permanent migrants to the ER provinces are not as over-represented in the extractive industries and construction as the long distance commuters. Instead, they look more like the existing residents of the ER provinces. The commuting process is more targeted at the high growth industries than either permanent migration or the existing labour force. The commuters to other provinces are dominated by those in public administration jobs, suggesting that many of them may be people who live in Quebec but cross over to Ottawa to work for the federal government.

Tables 1 and 2 show that there is substantial variation in the industrial composition (and, therefore, in the constructed rent variables) across our worker types. However, the number of observations at the bottom of the columns in each table indicate that the proportions of workers who either commute or migrate are small. Commuters to ER provinces are only 0.16% of residents in the non-ER provinces in 2000. Although this triples to 0.51%, the numbers are obviously small. Importantly, though, there is substantial variation across locations in this proportion. For Cape Breton, a relatively poor, ex-mining region in Nova Scotia, the percentage of all workers commuting to an ER province increases from 0.7% in 2000 to 6.5% in 2012. For young men, aged 22 to 34 living in Cape Breton, 1 in 8 of them were commuting to ER provinces in 2012. In contrast, for Toronto, the percentage of such commuters was 0.1% in 2000 and 0.2% in 2012. Thus, there is considerable cross-location variation in the changes in commuting proportions. In comparison, the percentage of workers in Cape Breton in 1999 who migrated to an ER province in 2000 was 0.50%, rising to 0.98% in 2012. The same numbers for Toronto are 0.06% in 2000 and 0.19% in 2012. Together, these imply that there is an increase in permanent movement but it is less salient than long distance commuting, especially in the Maritimes. Recall from the model that even small numbers of commuters may have substantial effects on non-ER province wages because workers in all sectors can point to the commuting option as a realistic part of their outside option. On the other hand, there were reasons to question the likely impact of permanent migration on bargained wages of the residents.

Table 3 shows the changes in the industry premia (ν_i's) between 2000 and 2012 computed with (first column) and without (second column) the 7 x 7 matrix of earnings/employment/self-employment positions in the previous two years. The estimates show substantial growth in industry premia in oil and gas, pipeline transportation, mining, and construction across these years. Including the 7x7 matrix generally implies smaller changes in industry premia, implying that there were some movements across earnings types that were associated with industrial changes in this period. We are concerned that those movements reflect selection on ability or education, which is why we estimate specifications including the 7x7 matrix.

5. Results

In Table 4, we present the results from our main specification, using all resident workers. Recall that we are pooling 4-year differences from three periods: 2000-2004, 2004-2008, and 2008-2012 and that our dependent variable refers to the mean log wages of residents, i.e., those who remain in their home, non-ER region and do not either commute or move. All standard errors are clustered at the economic region level. The first column contains OLS estimates of the specification that emerges from our theory with changes in the mean log wage in a (non-ER province) location x industry cell regressed on the change in the average rent in that location (ΔR_{ct}), the expected gain from moving permanently to an ER province (ΔX^P_{Act} - which incorporates both the probability of a worker from c moving to an ER province and the average rent in ER provinces of recent migrants from c), the expected gain from commuting to an ER province (ΔX^T_{Act} - which incorporates both the probability of a worker from c commuting to an ER province and the average rent in ER provinces of commuters from c), the same variables for non-ER province moves and commuting (ΔX^P_{Bct} and ΔX^T_{Bct}), and the change in the employment rate in c ($\Delta EmpR_{ct}$). The location specific rent and employment rate have positive and significant (at the 5% level) effects on the average wage in the city. These fit with the predictions from our theory: higher average rents in a local economy imply that workers in all sectors can bargain higher wages, and tighter labour markets as reflected in higher employment rates mean that relative bargaining power shifts toward workers. The effects of the rent variables related to the various types of migration/commuting options are not close to statistical significance with the exception of the commuting to ER provinces option, which is significant at the 10% level.

In the second column of Table 4, we present instrumental variables results in which we instrument for all the rent variables and the employment rate variable using the various types of Bartik instruments described earlier. The Sanderson-Windmeijer F-statistic values indicate that we do not face weak instrument problems with the possible exceptions of the ΔX^P_{Bct} (which has an associated p-value for the test statistic of 0.071) and $\Delta EmpR_{ct}$ (with an associated p-value of 0.045). The estimated effects of the outside options related to moving and commuting to non-ER provinces are swamped by their standard errors. We believe that this stems from the relative similarity of the industrial composition of employment for these movers relative to the residents. There is little clear extra advantage to be had from migrating to one of the other non-ER provinces in this period and, so, those options have little impact on negotiated wages. Given this, our preferred estimates do not include the non-ER province rent variables.

The estimated coefficient corresponding to the permanent move to ER provinces option is similarly poorly defined. Recall from our model that the option of moving permanently to an ER province will not affect wage setting in non-ER locations in equilibrium because the wage increases in the ER provinces will be matched with housing price increases. However, in the short run - if there is a period in which housing prices do not keep pace with wage increases - there could be an effect on non-ER province wages. Whether the permanent move option affects wages in our equation can, as a result, be seen as an indicator of whether our 4 year differences represent short or longer term adjustments. In fact, over different specifications we have estimated, the permanent move option effect has estimated sizes and even signs that switch wildly across specifications, but those estimates are never statistically significant and are almost always smaller than their associated standard errors. We take this as evidence that this effect is not well identified and interpret it is as implying that we are in a longer run situation in which the permanent move option is not relevant. For that reason, we also drop this variable in our remaining specifications.

The IV estimates omitting the non-ER province rent variables and the permanent ER province move rent variable are given in column 4 (with column 3 containing the OLS estimates for comparison). It is important to note at this point that in these more constrained regressions, we encounter a weak instrument issue with the employment rate. The p-value associated with the the Sanderson-Windmeijer F-statistic related to the employment rate is 0.16 in this case and is at least this size in other variants (defined by whether or not we include the 7 x 7 matrix, panel versus non-panel, and the geographic extent of the sample) of the regression we have run. Corresponding to this, the estimates of the employment rate effect are somewhat erratic when we instrument for the employment rate. In response, we proceed by instrumenting for ΔX^{T}_{Act} and ΔR_{ct} but not $\Delta EmpR_{ct}$. Following Stock and Watson(2011), we then interpret $\Delta EmpR_{ct}$ as what they call a 'control' variable, i.e., a variable whose coefficient picks up its own causal effect and any other unobserved factors that are correlated with it. That is, we use the employment rate variable to control for general demand related effects so we can focus on the estimates of the causal effects we care about: the wage-spillover effects from outside options. Given our earlier discussion, the identification requirement is then that start of period industrial composition in a city is not correlated with growth in city specific effects, where the latter are defined after conditioning out any such effects that are themselves correlated with changes in the city employment rate.

The estimate of the effect of the change in rent in the local economy is 0.96 and is statistically significantly different from zero at the 5% level. Green(2015) shows that this estimated effect has a direct relationship to a standard shift share calculation of the impact of a shift in industrial composition on the average wage in an economy. A shift in composition toward higher paying industries would alter the mean wage in a manner obtained by multiplying changes in industrial shares times industry premia and then summing across industries (the standard "between" effect). We can call that direct composition shift effect, B. Once we incorporate spill-over effects through bargaining, the total effect of the composition shift is B*(1 + β_1), where β_1 is the estimated coefficient on ΔR_{ct}. So, in our case, the spillover effects imply a doubling of the standard composition effect. This estimate is about a third of the corresponding estimates in Beaudry et al. (2012) for the US and Green (2015) for Canada. However, as we will see shortly, in estimates for the whole economy we get a ΔR_{ct} coefficient that is similar to those earlier studies.

The effect of the variable capturing the option of commuting to an ER province (ΔX^T_{Act}) also has the predicted positive effect and is statistically significant at the 5% level. Thus, an increase in the expected value of the commuting option generates increases in wages for the workers who do not take up that option. The employment rate again enters with an effect that has the predicted positive sign. Just as important for us, are the implications of the inclusion of the employment rate for the interpretations of the other coefficients. Effects of changes in the city level rent, ΔR_{ct}, might be interpreted as reflecting general increases in demand in the city and/or negative labour supply effects within most industries as workers move to the higher rent industries. However, in the absence of perfectly inelastic schedules, both of those mechanisms would be associated with changes in employment. The fact that we see wage effects of ΔR_{ct} while holding the employment rate constant fits with an interpretation of that effect as reflecting the wage bargaining spillovers emphasized in the model.[11] Similarly,when we control for the employment rate, the coefficient on ΔX^T_{Act} reflects bargaining effects and does not include effects resulting from reductions in local labour supply because of the increase in commuting.

In column 5, we present a specification in which we drop the employment rate variable. This means that the estimated coefficient on ΔX^T_{Act} reflects both the bargaining effect plus the effect of reducing labour supply as commuters shift their labour to an ER province destination and/or any demand effects resulting from the extra income the commuters spend in their home communities. The resulting estimate is about 25% larger than when including the employment rate. This fits with our claim that direct effects such as through supply shifts will be smaller than wage bargaining effects. Wage bargaining effects can be larger because workers in all sectors in the economy can refer to the commuting option in bargaining with their employer.

To provide perspective on the size of our estimated commuting option effects, in Table 5 we present actual and predicted changes in the mean log wage for various demographic groups and locations over the 2000 to 2012 period. As in our estimation, the numbers in this table correspond to residents - i.e., workers in the home, non-ER regions, excluding long

[11] Beaudry et al. (2012) argue further that the bargaining model implies that IV1 and IV2 should generate the same estimated effect if the identification conditions for the model hold. In our estimates, we find that we cannot reject the over-identifying restriction that estimates using only IV1 or IV2 are the same at any standard level of significance. When we estimate without including the permanent move to ER province option (the equivalent of column 4 in the table), we obtain an estimated coefficient on ΔR_{ct} of 0.99 when we use IV1 and 0.96 when we use IV2.

distance commuters and movers. We show the changes for two demographic groups: both genders, age 22 to 64 (our complete sample); and males, age 22 to 34. The latter group tends to have much higher probabilities of commuting and, therefore, are more likely to have their wages affected by the commuting option. To form the predictions in the second column of numbers, we use a combination of the estimated commuting effect option from column 4 of Table 4 and the observed value of ΔX^T_{Act} for the relevant group. For all seven non-ER provinces together, the mean real wage rises by 7% between 2000 and 2012 for the group containing both genders and all ages. In comparison, we predict a 0.9% increase, or 14% of the actual increase, based just on the change in the expected value of the commuting option. For young men, we predict a larger, 2.0% increase, because their probability of commuting increases more than the average. This over-explains what is essentially a zero real wage change for this group in our period.

These overall numbers conceal substantial variation across regions based on differences in the commuting probabilities. The Maritime provinces (Nova Scotia, New Brunswick, and PEI), which were an important source of commuting labour to the oil fields, have predicted effects that are approximately 10 times those for Ontario, where long distance commuters were much less prominent. This matches interestingly with figure 5 in which the time series pattern for the mean wage in the Maritimes mimics the pattern in the ER provinces to a much greater extent than does Ontario. Our results indicate that this is because of the stronger commuting ties between the Maritimes and the ER provinces. The largest commuter sending region was Cape Breton where approximately 1 in 8 young men commuted to ER provinces at the height of the boom. Given the depressed state of the local economy in Cape Breton related to the decline of coal jobs in the years before the oil boom , the 12% increase in the mean real wage for young men between 2000 and 2012 is striking. Our estimates indicate that the rise in the commuting option can completely account for this increase. In contrast, in Toronto, where less than 1% of young men commuted, the real wage for young men declined by 15% in this period and we estimate that changes in the value of the commuting option had essentially zero impact on the wages of the resident workers (i.e., those who did not commute). Thus, our results point to substantial wage spillover effects from the resource boom for those who did not take any direct role in the boom for some non-resource communities but not in others. It is worth reiterating that these estimated effects correspond to the bargaining power channel only. Once we allow for labour supply effects and commuting related demand effects by using the coefficient estimate from the column in which we do not control for the employment rate, these estimates are inflated by 25%.

In the last column of Table 4 we present the reduced form, regressing the wage on the oil price based instrument. Recall that this instrument is the percentage change in the oil price times a base year probability of workers from the location commuting to an ER province. The estimated coefficient can be interpreted as the total effect of the oil price increase, operating through the channels we have so far discussed (shifts in bargaining power and labour supply due to the expansion and value of the commuting option) but it will also reflect effects through shifting the industrial composition in the non-ER province location. The latter could arise because of Dutch disease effects arising because of increased wages and exchange rates affecting local, high wage industries. For the entire non-ER region and our full demographic sample, the reduced form coefficient combined with the actual changes in the instrument value imply a 0.5% increase in the real mean wage for residents. This is less than the 0.9% increase we get from the more structural estimates, which would fit with negative Dutch Disease type effects arising because of reductions in job creation by local businesses facing higher wage costs.

In Table 6, we present two further sets of estimates. In the first, we do not use the 7x7 matrices in the estimations in which we obtain the industry rents and our dependent variable. We view the 7x7 matrices as useful for allowing us to argue that what we are using is, at least roughly, industry rents rather than differences in education and other skills. It is the former not the latter which are the relevant source of variation in our model.[12] The results including the 7x7 matrices are our preferred estimates but we present the estimates without those controls out of concern that we may be throwing out too much of the relevant variation. The first three columns of 6 provide estimates from our main specification (without the permanent migration options and the commuting option to non-ER provinces) estimates by OLS and by IV with and without the employment rate control. These results show the same broad patterns as those in Table 4. In particular, the location rent variable and the commuting to ER provinces effects are positive and generally statistically significant in the IV specifications. The employment rate continues to have a positive effect on wages and removing it from the specification leads to an increase in the coefficient on ΔX^T_{Act}. The estimated effects for ΔR_{ct} are of similar size to those in Table 4 but the coefficients on ΔX^T_{Act} are considerably smaller. However, the values for ΔX^T_{Act} are larger when the industry premia are obtained when not controlling for prior earnings history. For all non-ER regions combined, the increase in ΔX^T_{Act} from 2000 to 2012 combined with the estimate from the specification including the employment rate implies a 1.0% increase in the real mean wage compared to 0.9% when we use the 7x7 past earnings matrix. Thus, the results from the two approaches are similar in size.

To this point, we have presented specifications in which we use instrumental variables to address issues of endogeneity of our right hand side variables; trying to make sure that our rent variables, in particular, are not actually picking up shifts in demand conditions in a location. Without this approach, we might expect that we will capture a greater propensity to migrate out of regions currently having economic troubles and, indeed, our instrumental variables estimates of the impact of the commuting to ER provinces option are larger than our OLS estimates of that effect. But these instruments do not address

[12]Industry differentials based on skill differences would not have wage bargaining spillover effects of the type described in our model because, for example, a high school drop-out worker could not point to increases in average wages in town that stem from the arrival of a highly educated sector as a relevant outside option. That means that to the extent that we do not strip skill differences out of our initial stage industry differential estimates, we will be mixing relevant variation (from rents) with variation that we would expect to have zero effect (from skills). To the extent this is true, less complete control for skills should imply smaller estimated rent variable effects. At the same time, we are working with differences in industry differentials as our rent measure. If the educational composition of industries and the returns to education by industry do not vary with time then that differencing could remove any of these concerns.

the other potential endogeneity concern: the selectivity of migrants. Recall that our wage regressions are estimated using the population of people who remain resident in (and do not commute from) the non-ER provinces. But if the people who migrate and commute are, for example, less good workers then we would expect an increase in the mean observed wage with increases in the value of ΔX^T_{Act} simply because of the composition change of the non-commuting workers.[13] We attempt to address this by using our panel sample of workers. This consists of the set of workers who work in non-ER provinces for all 4 of our data years in a given 4 year difference. As stated earlier, while this is undoubtedly a select group, it is a consistent group across years and thus does not change in terms of its composition of observable or unobservable characteristics.

The results using our consistent panel of workers are presented in columns 4 through 6 of Table 6, where we again use our main, more parsimonious specification and use the 7x7 matrix of past earnings. The estimated effect for ΔR_{ct} and ΔX^T_{Act} are again statistically significant and positive but are smaller in size than the matching estimates not using the consistent group presented in Table 4. However, once we take account of the actual changes in these variables for the consistent sample of people who never moved or commuted, the implied effect of the commuting option (when controlling for the employment rate) is a 0.7% increase, which is smaller than the 1.0% increase for the pooled sample when we do not control for the 7x7 earnings matrix and 0.9% when we do. Thus, these results suggest that the selection of migrants and commuters generated a slight upward pressure on the estimated spillover effects of the commuting option onto non-commuters. However, our main conclusions and the size of our estimated effects are not substantially altered by accounting for selection.

6. Accounting for the Role of the Resource Boom in Canadian Wage Changes

In this section, we make use of the estimates presented to this point along with other calculations to obtain an estimate of the impact of the resource boom on the increase in the Canadian mean real wage between 2000 and 2012. In order to connect this exercise, based on earnings data, back to our initial wage movement figures, we first investigate the relationship between earnings and wage movements in our period. To that end, in Table 7, we present log wage and earnings values for Canada and the US in Canadian Census years (2000, 2005, and 2010). In the first two columns, we present the real mean log hourly wages for Canada and the US in those years from Figure 1 (i.e., based on LFS and CPS data). In the third and fourth columns we use Canadian Census data for 2000 and 2005, data from the 2011 Canadian National Household Survey (which refers to the year 2010), data from the 2000 US Census, and data from the 2005 and 2010 American Community Surveys to construct real mean weekly earnings.[14] The fifth and sixth columns contain mean log annual earnings values from the Census type sources, and the eighth column contains mean log annual earnings from our tax data. In all the matched pairs of columns, one can again

[13] Recall that ΔX^T_{Act} is partly a function of the proportion of people from the location who commute.
[14] In all cases, we restrict our attention to 20 to 54 year olds.

see the much better performance of wages in Canada relative to the US. As we move to earnings concepts that include weekly hours and then, in addition, weeks worked per year, the US performance worsens - from a 1% drop in the hourly wage between 2000 and 2010 to an 8% drop in the weekly wage to an 11% drop in annual earnings. In contrast, Canada experienced a 7% gain in real hourly wages, a 9% gain in real weekly wages, and either a 6% or 5% gain in real annual earnings depending on whether one uses Census data or tax data. The relative constancy of changes across the different measures reveals, again, that Canada had better outcomes both in terms of wages and in hours of work after 2000. In addition, that constancy implies that we can reasonably take our estimated effects from the tax based annual earnings data and apply them to the wage data since hours and weeks worked seem to have changed little across this period in Canada on average.

We are interested in how much of the observed 7.6% increase in mean real hourly wages between 2000 and 2012 can be attributed to the resource boom. To calculate this, we construct a series of counterfactual exercises. The first of these corresponds to the direct effect of changes in the extent of employment in the ER sector in ER provinces and changes in the wage premium on changes in the overall mean real wage for the country as a whole. Given the increase in the proportion of workers in this sector in ER provinces from 0.059 to 0.092 and the increase in the sectoral wage premium from 0.27 to 0.34 calculated in (Fortin and Lemieux (2015)) from LFS data , the direct effect of the resource boom was a 1.5% increase in the mean hourly wage in the ER provinces. Combined with the fact that the ER provinces account for 16% of total Canadian employment, the implied effect on the overall mean wage for Canada is an increase of 0.24%. This, again, highlights that focusing just on those in or joining the resource sector results in small overall effects.

In our next step, we take account of spillover effects on wages within the ER provinces. To do this, we adopt the specification from (Beaudry et al., 2012) and (Fortin and Lemieux, 2015) in which we regress mean log earnings from our tax data in industry - economic region cells on ΔR_{ct} and the change in the employment rate. In contrast to our earlier estimates, we now include data from all 73 economic regions in the country. We also allow for the possibility that the resource sector is particularly salient. The argument in our model is that workers in all sectors can point to the arrival or expansion of a high rent sector as part of their outside option in bargaining with their employers. High rent sectors that get a lot of attention in the press and in public conversation should be particularly useful in that bargaining, and the oil sector certainly fit that description during the boom. To allow for resource sector wage rents to have a particularly strong effect, we add in an extra regressor which corresponds to the resource sector's component of ΔR_{ct}, i.e., the change in the proportion in the resource sector times its wage premium. We instrument for this variable using the interaction of the proportion of employment in the ER sector in the region in 2000 and the change in the price of oil during the four year period over which the difference is taken. We present the results from this specification with and without the ER specific rent share variable in Table 8. As before, we do not instrument for the change in the employment rate. The estimated effect of ΔR_{ct} is approximately 3 when not including the ER rent share variable, which is very similar to what (Beaudry et al., 2012) obtain for the US and (Green, 2015) obtains using Census data for Canada. When we include the ER rent share variable the coefficient on that variable is very large and the coefficient on ΔR_{ct} drops to 2, implying that the resource rents played a large role in wage bargaining in Canada in this period.

We obtain our estimate of the effect of the resource boom on the mean wage in the resource sector by multiplying one plus the sum of the ΔR_{ct} coefficient and the coefficient on the ER rent share variable times the change in the ER rent share variable between 2000 and 2012.[15] The latter variable is just the sum of the direct employment share and rent increase effects calculated in the first step (i.e., 1.5%). The total effect with the ER provinces is an implied 0.16 log point increase in mean wages. This fits with arguments about local spillovers from resource booms in Feyrer et al. (2017), Marchand (2015), and Fortin and Lemieux (2015) among others. To get the effect on the overall Canadian wage, we multiply this by the proportion of employment in the ER provinces, yielding an overall implied increase of 2.6%. Note that this includes both the direct effects for resource workers and the spillovers to wages of other workers in ER provinces.

The next step is to add in the bargaining spillover effects related to the long distance commuting option in the non-ER provinces. As we have seen, this amounts to a 0.9% increase in the mean wage in those provinces, or a 0.78% increase in the mean wage for Canada as a whole. When we allow for labour supply and commuter spending effects in those regions, this is inflated to 1.0%.

Finally, we allow for spillovers in the form of derived demand for goods produced in other parts of Canada that are sold to the ER sector. For example, 12% of sales from machinery manufacturing was sold to the ER sector in Canada in 2010. Importantly for our discussion, machinery manufacturing is a high wage rent industry, paying a wage premium of 15% relative to the average in 2000. Thus, an oil boom induced expansion in machinery manufacturing in a city would raise the value of R_{ct} and, with it, the outside option for workers bargaining in all sectors. We carried out a rudimentary exercise to approximate the role of this channel in wage changes in the non-ER provinces. In particular, we used input-output tables for 2010 to calculate the proportion of total sales sold to Canadian ER sector firms in each NAICS 3 digit industry. We then assumed that demand by the ER sector grew proportionally to the rise in the oil price (i.e., by 0.3 log points from 2000 to 2012). Using the ER sales proportions and the growth rate, we predicted how much smaller (in terms of employment) each industry in each of our non-ER regions would have been if there had been no ER sector growth. From this, we calculated a counterfactual version of the rent variables, R_{c2010} and compared the actual change, ΔR_{ct} between 2000 and 2012 with the counterfactual change with no resource boom. Using this in combination with our coefficient on ΔR_{ct}, we calculate that the resource boom raised mean wages in the non-ER provinces by approximately 0.1%. Thus, while there was some effect on production in non-ER provinces, the share of sales to the ER sector is simply too small to have much effect as an outside option for other workers. Returning to our machinery manufacturing example, the Windsor region in western Ontario has one of the largest machinery manufacturing sectors but even

[15]We add together the coefficient on ΔR_{ct} and the coefficient on the ER rent share variable since the latter is part of the former and thus changes in the rent share interact with both coefficients. We add one for reasons described earlier and discussed in more detail in (Green, 2015).

there it constitutes only 2.7% of total employment. Considering that only about 12% of that 2.7% is being sold to the ER sector directly, it is easy to see that the direct effect numbers get small very quickly. While there was talk in Canada during the boom of spillovers of this type, our calculations - even with the inflation that comes with bargaining spillovers - suggest that this was not an important component of wage determination.

We present the steps in our decomposition all together in Table 9. Out of the total 0.076 increase in the mean log wage for Canada between 2000 and 2012, our calculations imply that 49% can be attributed to the resource boom. The great majority of this stems from spillovers within the ER provinces themselves and spillovers to other communities through the long distance commuting option. Spillovers through induced demand of the ER sector for goods produced in other provinces play a very minor role.

7. US Results

In this section we take our long distance commuting and wage spillovers specification to US data. In a first exercise, we attempt to recreate our Canadian estimates of the impact of the oil boom on wages in non-oil states. The importance of long distance commuters in resource extraction in the US is well known. For example, Carrington (1996) states that virtually all the skilled pipefitters and engineers and many of the lower skilled workers who worked on the Alaskan pipeline were from the lower 48 states and worked in Alaska on contract. We are interested in the effects of these commuters on wages in their home states. We investigate this using the 2000 US Census and the American Community Survey(ACS) grouped into 3 pairs of years: 2005-6; 2009-10; and 2014-15. We group years in order to get sufficient sample size. The public use versions of the Census and ACS contain individual States of residence but also the state where the person works. We treat people who live in one state but work in another as long distance commuters and the people who live and work in the same state what we have called 'residents'. Given the results in the Canadian case, we do not attempt to identify permanent movers. We are interested in long distance commuters to ER boom states, which we define as Alaska, North Dakota, New Mexico, Oklahoma, Louisiana, Texas, West Virginia, and Wyoming based on their proportions of employment engaged in the ER sector. We use the Consistent Public Use Microdata Areas (CPUMA) developed by IPUMS as our 'cities', and aggregate the 1990 industry classification into 45 consistent industries. We then run our parsimonious specification (the one that does not include permanent moves or long distance commuting to non-ER states) with the units of observation being city by industry cells in the non-ER states. The dependent variable is the change in the average log weekly wage in the given city by industry cell. We construct ΔR_{ct}, ΔX^T_{Act} and the change in the city employment rate along with the instruments for ΔR_{ct}, ΔX^T_{Act} in the same way as in the Canadian data. [16] More details on data construction are provided in Appendix C.

[16]Thus, X^T_{Act} is constructed as the proportion of people living in city c who are recorded as working in an ER state times the expected rent in the ER state for long distance commuters. The latter is constructed using the proportions of long distance commuters in each industry in ER states and the national level industrial premia.

In the first four columns of Table 10, we present results from our US estimation both with and without the commuting option variable and estimated either by OLS or IV. The statistics associated with the first stages for all our endogenous regressors indicate that we do not face any weak instrument problems. Standard errors are clustered at the CPUMA level. Our preferred specification is the IV estimation in the fourth column, where we observe that the coefficient on ΔR_{ct}, showing the effects of spillovers from industrial composition in the local area, is over 2. This is similar in magnitude to what Beaudry et al. (2012) find using US Census data over a longer period. The coefficient on the commuting option is remarkably similar to what we observed in Table 4 for Canada and is statistically significantly different from zero at any standard level of significance. Thus, long distance commuting to resource rich states has bargaining effects on wages for non-movers in the commuters' home states just as it does across provinces in Canada. The key difference between the two countries is the salience of the resource boom in the two countries. While the proportion of workers in the states sending the most commuters (Mississippi at 1.6% and Arkansas at 1.3% of their workforces in 2000) are comparable to the highest sending province in Canada (Nova Scotia), the importance for the country as a whole is less. As a result, the estimated coefficient combined with the actual changes in the expected value of the ER commuting option only amounts to a 0.15% increase in wages in non-ER states as a whole compared to a 0.9% increase for the non-ER provinces in Canada.

The second four columns of the table contain a parallel attempt to examine the effect of the post-2000 US housing boom on wages in areas not as directly affected by that boom. We work with the same Census and ACS data as for the US oil boom exercise and assign workers to 'cities' by using state and PUMA codes, using the 1999 definition of SMSA.[17] For workers working outside their PUMA of residence, we use the variable `pwpuma` to assign workers to the city in which they work. Workers are coded as 'commuters' if they report working outside their city of residence using the variable. Since some CPUMAs contain more than one cities, we require a worker to be working both outside his/her city of residence and CPUMA of residence to be labelled a commuter. In order to identify housing boom cities, we work with housing price data from Zillow at the county level. We use crosswalks provided by MAPLE project at the Missouri Data Centre to assign counties to pumas, and then pumas to cities. We then calculate housing price growth at the city level as the log change in house prices between 2000 and 2006. We classify a city as a 'boom' city if the city's house price growth is above the 75 percentile. We then classify commuters who commute to a boom city using this definition. We again estimate our specification using data from non-boom cities and construct the commuting option in the same way as in the oil boom exercise.

The results from the housing boom specification are remarkably similar to those from the oil boom specification. The effects of the own-city rent variable (R_{ct}) and the employment rate variable are the same size in the two exercises. The commuting option variable again

[17]Since the coding of PUMAs change in 2011, we construct a cross walk that maps the new PUMA codes into the 2000 PUMAS using codes provided by the MAPLE project at the Missouri Data Centre (http://mcdc.missouri.edu/websas/geocorr14.html).

enters with a positive and statistically significant effect and in our preferred specification in column 8 takes a value that is 50% larger than the oil boom commuting option variable in column 4. In contrast to the ER commuting option, the housing boom option is much more important in a variety of cities across the US, with the largest proportions appearing to be related to relatively short distance commutes. For example, two of the highest proportion commuters are Riverside, California and Ventura, California, both with over 20% of their workforce commuting in 2000 and both bedroom communities of LA (one of the housing boom cities). However, the increases in the proportion commuting during the housing boom were sufficiently small that the total implied impact on wages in non-housing boom communities between 2000 and 2006 was an increase of 0.12%. Thus, for the commuting option to have an important impact on wages, there must be a significant swelling of the ranks of commuters combined with an increase in the value of the option. This was the case for the resource boom in Canada but apparently not for the resource and housing booms in the US.

8. Conclusion

We examine the impact of Canada's resource boom on wages across the economy, asking whether the resource boom effects can explain the substantial differences in the evolution of the wage structure in Canada and the US after 2000. In particular, we focus on whether wages in non-resource local economies across Canada were affected by the resource boom through migration and long-distance commuting links. Our examination is based on administrative data which we can use to identify permanent migrants from local economies in non-resource intensive provinces to the resource intensive provinces. We can also identify people who are undertaking long-distance commuting to the extractive resource intensive provinces: people who file their taxes in a community outside an extractive resource sector but get their main earnings from a firm in an extractive resource sector. We find that locations with higher proportions of long distance commuters have larger increases in wages. For Cape Breton, a low wage economy on Canada's east coast, we estimate that the increase in the value of the commuting option because of the onset of the oil boom implies a 13% increase in the local average wage. We explain this as arising primarily from bargaining effects: non-commuters in Cape Breton can threaten to start commuting to the oil provinces as their cousins are doing and, using that threat, bargain a better wage. These same kinds of bargaining threats were happening within the extractive resource (ER) intensive provinces on a larger scale. We present estimates similar to those in Fortin and Lemieux (2015) indicating that the effects of the resource boom on wages in other sectors of the ER province economies was large. Combining these various effects, we can account for 49% of the increase in mean real wages in Canada between 2000 and 2012 - a number that leaves out other effects such as the expansion of the public sector as a result of the increased taxes related to the boom. In comparison, we find very similar size effects of the option of commuting to extractive resource states on wages in non ER states relative to our Canadian estimates. However, the option of commuting to work in the oil fields and mines during the resource boom was not as salient in the US as it was in Canada and the overall impact on US wages in non-ER locations was about 1 sixth of our estimated impacts in Canada.

The main implication from these results is that long distance commuting, which is a common feature of resource development, can serve to spread the effect of a resource boom over a much wider geographic region than previously suspected. Indeed, with the advent of low cost air fares, there is no real geographic restriction on this channel. For both Canada and the US, the commuter sending communities tend to be lower income locations where prior resource or manufacturing operations have declined. Cape Breton, for example, is an ex-coal mining region with perennially high unemployment. Approximately, 1 in 8 young men were making the trek to Alberta at the height of the boom, facilitated by frequent direct flights from Halifax (the nearest city to Cape Breton) to Fort McMurray (a boom town in Alberta's north that is near the oil fields). It is worth noting that with the end of the boom, there are no longer direct flights between Halifax and Fort McMurray. We argue that the expansion of these types of commuting options can have substantial effects on the wages of those who do not take up the option because of bargaining effects: the frequent direct flights and examples of many others taking up the commuting option strengthens the hands of workers in all sectors in wage bargaining. Since all workers in the local economy can point to the option even if they do not take it up, the effects can be widespread. Of course, these are only the wage impacts of the commuting. The dislocation of families as the men worked away from home for weeks at a time has potentially substantial impacts on the families, the men themselves and both the sending and receiving communities (Bartik et al., 2017). Those costs would need to be set against the wage benefits in a full accounting of the way the resource boom affected both non-resource and resource communities.

One implication of the finding that long distance commuting can spread the effects of booms is that previous estimates of the spillover effects from resource booms are mismeasured. The standard approach to estimating these effects is to compare wage and employment outcomes in resource boom locations to those in non-boom locations. Even papers that attempt to measure spillovers over wider regions focus on outcomes in areas within specified areas centred on the boom locations, comparing them to outcomes in locations that are farther away. In such an approach, Cape Breton would be treated as part of the 'control' group of locations rather than as one affected by the boom, leading to an under-estimate of the effect of the resource boom.

Our results indicate that the Canadian economy as a whole was strongly affected by the resource boom in the 2000s. This had positive implications in the increase in wages across the country and in the fact that the 2008 recession was not nearly as deep in Canada as it was in the US. On the other side of the ledger, the rise in wages in communities that did not have resource driven demand increases made labour more expensive in those communities and would be expected to cause a decrease in job creation by firms there (Beaudry et al., 2014). This combined with the rise in the value of the Canadian dollar associated with the boom could have implied Dutch disease effects in the non-resource regions and sectors of the economy. We have not pursued those implications further in this paper. However, it is worth noting that while any such Dutch disease effects may point toward concerns about the impact of trade with China and other economies that have been a prominent part of US discussions (e.g. Autor et al., 2016), the positive effects of the resource boom imply that expansions of trade had direct, positive economic benefits for Canada in this period. At the same time, the fact that those benefits came largely through the extraction of some of the world's dirtiest oil implies tough policy choices for Canada as it faces the dual challenges of meeting its climate accord promises and maintaining a strong economy.

The importance of the resource boom for Canadian wages means that most of the parting of ways in the wage structures between Canada and the US after 2000 can be explained by the relatively greater importance of the resource sector in the Canadian economy. This, in turn, means that comparisons of outcomes between Canada and the US in this period are likely not very useful for trying to uncover the impact of technological change on economies. While the technological factors that have been emphasized in the US are surely having an impact in Canada, their effects appear to have been swamped by resource sector expansion in the 2000s.

Acemoglu, D., Autor, D., June 2010. Skills, tasks and technologies: Implications for employment and earnings. NBER Working Paper 16082, National Bureau of Economic Research.
URL http://www.nber.org/papers/w16082

Allcott, H., Keniston, D., September 2014. Dutch disease or agglomeration? the local economic effects of natural resource booms in modern america. Working Paper 20508, National Bureau of Economic Research.
URL http://www.nber.org/papers/w20508

Antonczyk, D., DeLeire, T., Fitzenberger, B., Mar. 2010. Polarization and rising wage inequality: Comparing the u.s. and germany. IZA Discussion Papers 4842, Institute for the Study of Labor (IZA).
URL http://ideas.repec.org/p/iza/izadps/dp4842.html

Autor, D. H., Dorn, D., August 2010. The growth of low skill service jobs and the polarization of the u.s. labor market. Mit working paper, M.I.T.
URL http://econ-www.mit.edu/files/1474

Autor, D. H., Dorn, D., Hanson, G. H., October 2016. The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade. Annual Review of Economics 8 (1), 205–240.
URL https://ideas.repec.org/a/anr/reveco/v8y2016p205-240.html

Bartik, A. W., Currie, J., Greenstone, M., Knittel, C. R., January 2017. The local economic and welfare consequences of hydraulic fracturing. Working Paper 23060, National Bureau of Economic Research.
URL http://www.nber.org/papers/w23060

Beaudry, P., Green, D. A., Sand, B. M., May 2012. Does industrial composition matter for wages? an empirical evaluation based on search and bargaining theory. Econometrica 80 (3), 1063–1104.

Beaudry, P., Green, D. A., Sand, B. M., 2014. Spatial equilibrium with unemployment and wage bargaining: Theory and estimation. Journal of Urban Economics 79, 2–19.

Black, D., McKinnish, T., Sanders, S., 04 2005. The Economic Impact Of The Coal Boom And Bust. Economic Journal 115 (503), 449–476.
URL https://ideas.repec.org/a/ecj/econjl/v115y2005i503p449-476.html

Card, D., Heining, J., Kline, P., 2013. Workplace heterogeneity and the rise of west german inequality. Quarterly Journal of Economics 128 (3), 967–1015.

Carrington, W. J., 1996. The alaskan labor market during the pipeline era. Journal of Political Economy 104 (1), 186–218.
URL http://www.jstor.org/stable/2138964

Charles, K. K., Hurst, E., Notowidigdo, M. J., 2016. The masking of the decline in manufacturing employment by the housing bubble. Journal of Economic Perspectives 30 (2), 179–200.

Feyrer, J., Mansur, E. T., Sacerdote, B., April 2017. Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution. American Economic Review 107 (4), 1313–1334.
URL https://ideas.repec.org/a/aea/aecrev/v107y2017i4p1313-34.html

Fortin, N., Lemieux, T., 2015. Changes in wage inequality in canada: An interprovincial perspective. Canadian Journal of Economics 48 (2), 682–713.

Green, D. A., 2015. Chasing after 'good jobs'. do they exist and does it matter if they do? Canadian Journal of Economics 48 (2), 612–646.

Green, D. A., Sand, B. M., 2015. Has the canadian labour market polarized? Canadian Journal of Economics 48 (4), 1215–1265.

Jacobsen, G. D., Parker, D. P., 2016. The economic aftermath of resource booms: Evidence from boomtowns in the american west. The Economic Journal 126 (593), 1092–1128.
URL http://dx.doi.org/10.1111/ecoj.12173

Kline, P., 2008. Understanding sectoral labor market dynamics: An equilibrium analysis of the oil and gas field services industry. Tech. rep., UC Berkeley Economics.

Marchand, J., 2012. Local labor market impacts of energy boom-bust-boom in Western Canada. Journal of Urban Economics 71 (1), 165–174.
URL https://ideas.repec.org/a/eee/juecon/v71y2012i1p165-174.html

Marchand, J., 2015. The distributional impacts of an energy boom in western canada. Canadian Journal of Economics 48 (2), 714–735.

Michaels, G., March 2011. The Long Term Consequences of Resource-Based Specialisation. Economic Journal 121 (551), 31–57.

Morissette, R., Chan, P. C. W., Lu, Y., 2015. Wages, Youth Employment, and School Enrollment: Recent Evidence from Increases in World Oil Prices. Journal of Human Resources 50 (1), 222–253.
URL https://ideas.repec.org/a/uwp/jhriss/v50y2015i1p222-253.html

Slater, J., October 2016. Americans on the edge. The Globe and Mail.

Tschopp, J., Sep. 2015. Wage Formation: Towards Isolating Search and Bargaining Effects from the Marginal Product. Working Papers 055, Ryerson University, Department of Economics.
URL https://ideas.repec.org/p/rye/wpaper/wp055.html

Weber, J. G., 2012. The effects of a natural gas boom on employment and income in colorado, texas, and wyoming. Energy Economics 34 (5), 1580 – 1588.
URL http://www.sciencedirect.com/science/article/pii/S0140988311002878

Table 1: Industrial Composition for Residents, Commuters, and Migrants, 2000

	Residents	Connections with ER provs Commuters	Connections with ER provs Migrants	Connections with Other provs Commuters	Connections with Other provs Migrants	Workers residing in ER provinces
			percent			
Agriculture	1.7	4.1	2.5	1.8	1.4	2.2
Mining, Oil	0.6	8.5	4.1	0.8	0.6	5.0
Utilities	0.9	0.4	0.4	0.2	0.3	0.9
Construction	4.6	18.2	9.3	6.0	3.2	6.7
Manufacturing	18.5	5.9	10.6	9.1	12.3	9.2
Wholesale	5.4	3.4	4.6	4.6	4.5	5.3
Retail	9.8	5.0	9.9	4.3	8.5	10.4
Transportation	4.5	7.7	5.7	5.5	4.9	5.2
Culture	2.5	1.3	1.9	2.6	3.3	2.4
Finance	4.9	1.9	3.3	2.6	4.6	3.6
Real Estate	1.3	1.5	1.4	0.8	1.2	1.6
Professional	4.9	6.2	5.5	4.3	7.0	5.3
Management	0.5	0.9	0.5	0.5	0.5	0.6
Waste	4.0	5.1	5.4	5.9	6.3	3.4
Education	7.9	4.5	6.2	6.6	8.2	8.7
Health	9.8	2.7	6.5	4.2	7.7	9.2
Arts	1.2	1.7	1.6	1.5	1.7	1.1
Accommodation	4.8	9.8	9.1	3.8	6.6	5.3
Services	3.9	3.0	3.8	3.2	3.7	4.4
Public Admin	8.3	8.4	7.9	31.8	13.8	9.9
No. of obs	9,028,875	14,406	17,034	106,208	64,744	1,697,408

Note: Numbers may not add to 100.0 due to rounding. Except for migrants, the sample consists of paid workers who resided in non-oil-producing provinces in 2000.

Table 2: Industrial Composition for Residents, Commuters, and Migrants, 2012

	Residents	Connections with ER provs Commuters	Connections with ER provs Migrants	Connections with Other provs Commuters	Connections with Other provs Migrants	Workers residing in ER provinces
			percent			
Agriculture	1.4	1.9	1.2	1.4	1.2	1.4
Mining, Oil	0.6	16.3	6.6	1.9	1.0	6.5
Utilities	0.8	0.7	0.7	0.2	0.3	1.2
Construction	6.0	33.7	12.1	5.7	3.8	9.3
Manufacturing	12.0	2.9	6.9	6.4	6.3	6.8
Wholesale	5.0	2.9	3.9	5.7	3.6	4.7
Retail	10.4	2.6	9.7	3.9	9.5	9.8
Transportation	4.4	7.2	5.2	4.9	3.7	4.8
Culture	2.3	0.5	1.4	1.9	2.8	1.6
Finance	5.1	0.8	2.5	2.5	4.1	3.2
Real Estate	1.6	2.0	1.9	0.8	1.3	1.8
Professional	5.9	6.6	7.1	5.8	7.7	6.6
Management	0.8	0.4	0.6	0.6	0.6	0.6
Waste	4.9	5.9	6.8	5.2	6.1	3.9
Education	8.3	2.3	5.0	7.7	8.6	7.8
Health	11.8	1.6	6.7	5.2	9.0	8.1
Arts	1.2	0.8	1.2	1.3	1.6	1.1
Accommodation	5.2	5.6	8.3	3.6	8.6	5.0
Services	3.9	2.4	3.9	2.9	3.9	4.3
Public Admin	8.6	3.0	8.4	32.6	16.5	11.6
No. of obs.	10,053,946	52,611	33,837	122,076	54,420	2,116,328

Note: Numbers may not add to 100.0 due to rounding. Except for migrants, the sample consists of paid workers who resided in non-oil-producing provinces in 2012.

Table 3: Growth in adjusted log earnings, 2000-2012, Selected Industries

Version 1	Version 2	Industry
0.17	0.24	oil and gas
0.14	0.23	Mining
0.10	0.25	Utlities
0.02	0.14	Bldg Constr
0.07	0.20	Heavy Const
-0.05	-0.11	Textile prods
-0.07	-0.18	Wood prod
0.03	0.02	Petr prod
-0.05	-0.09	Chem manf
-0.17	-0.14	comp manf
-0.08	-0.16	trans manf
0.15	0.25	gas distrib
-0.10	-0.12	furniture store
-0.12	-0.09	elec stores
-0.10	-0.18	air transp
0.27	0.30	pipelines
-0.10	-0.09	financial
-0.06	-0.07	education
-0.04	-0.05	health
-0.12	-0.13	recreation
-0.04	-0.04	accomod.
-0.03	-0.02	food retail
-0.14	0.11	public admin

Note: Changes measured relative to base group: crop production. Version 1: log wage regressions include the 7x7 matrices corresponding to earnings quintiles, nonemployment, and self-employment as well as full interactions of gender, immigrant status, and a quadratic in age.. Version 2 based on log wage regressions with full interactions of gender, immigrant status, and a quadratic in age but without the 7x7 matrices.

Table 4: Estimation Results, Canada, Non-Panel Based, Including 7x7 Matrices

	OLS	IV	OLS	IV	IV	OLS
Variables	(1)	(2)	(3)	(4)	(5)	(6)
ΔR_{ct}	1.00*	0.79	1.00*	0.96*	1.12*	
	(0.32)	(0.51)	(0.31)	(0.32)	(0.37)	
ΔX^P_{Act}	-7.37	5.78	-	-	-	
	(8.06)	(16.51)				
ΔX^T_{Act}	2.84	4.61	2.32	7.78*	9.68*	
	(1.46)	(2.88)	(1.30)	(2.03)	(2.52)	
ΔX^P_{Bct}	-4.00	2.61	-	-	-	
	(3.24)	(14.83)				
ΔX^T_{Bct}	0.21	-0.17	-	-	-	
	(0.40)	(3.60)				
$\Delta EmpR_{ct}$	0.41*	0.92*	0.42*	0.42*	-	
	(0.08)	(0.28)	(0.08)	(0.08)		
oilTA						0.034*
						(0.007)
Industry	yes	yes	yes	yes	yes	
Year	yes	yes	yes	yes	yes	
Observations	14,247	14,247	14,247	14,247	14,247	
R^2	0.41		0.41			0.40
Instrument set		IV1s - IV2s		IV1, IV2	IV1, IV2	
		IV3 oilTA				
		oilPA		oilTA	oilTA	
Sanderson-Windmeijer		Multivariate	F tests			
ΔR_{ct}		19.52		386.94	336.11	
ΔX^P_{Act}		4.63		-	-	
ΔX^T_{Act}		10.91		36.17	36.50	
ΔX^P_{Bct}		1.95		-	-	
ΔX^T_{Bct}		2.54		-	-	
$\Delta EmpR_{ct}$		2.17		-	-	
P-values for S-W tests						
ΔR_{ct}		0.000		0.000	0.000	
ΔX^P_{Act}		0.002		-	-	
ΔX^T_{Act}		0.000		0.000	0.000	
ΔX^P_{Bct}		0.071		-	-	
ΔX^T_{Bct}		0.020		-	-	
$\Delta EmpR_{ct}$		0.045		-	-	

Note: Standard errors, in parentheses, are clustered at the economic region level. The asterisk (*) denotes significance at the 5% level. All models are estimated using 55 economic region by 102 industry cells in four year differences.
ΔR_{ct} is average rent in the economic region. ΔX^P_{Act} is the change in: the probability a person from location c moves permanently to an ER province times the average rent of new migrants in ER provinces. ΔX^T_{Act} is the change in: the probability a person from location c commutes to an ER province times the average rent of commuters in ER provinces. ΔX^P_{Bct} and ΔX^T_{Bct} are defined analogously for non-ER province destinations. $\Delta EmpR_{ct}$ is the change in the economic region level employment rate.

Table 5: Actual and Predicted Wage Changes, 2000 to 2012

	Actual	Predicted	Proportion
Non-ER Provs			
Both Genders, 22-64	0.066	0.009	0.14
Males, 22-34	-0.011	0.021	-1.86
Maritimes			
Both Genders, 22-64	0.17	0.041	0.25
Males, 22-34	0.13	0.079	0.61
Cape Breton			
Both Genders, 22-64	0.19	0.13	0.70
Males, 22-34	0.12	0.25	2.07
Ontario			
Both Genders, 22-64	0.026	0.004	0.15
Males, 22-34	-0.11	0.010	-0.093
Toronto			
Both Genders, 22-64	-0.012	0.002	-0.17
Males, 22-34	-0.15	0.0058	-0.039
BC			
Both Genders, 22-64	0.024	0.024	1.00
Males, 22-34	0.014	0.067	4.81

Note: The actual column shows the proportional changes in the mean wage for the specified location and demographic group for the period from 2000 to 2012. The predicted column shows the changed predicted based on the commuting option coefficient in column 4 in Table 4 combined with the value for ΔX^T_{Act} for the period 2002 to 2012.

Table 6: Estimation Results, Alternative Specifications

Variables	Non-panel Without 7x7 Matrix OLS (1)	IV (2)	IV (3)	Panel With 7x7 Matrix OLS (4)	IV (5)	IV (6)
ΔR_{ct}	0.82*	1.01*	1.09*	1.22*	1.23*	1.21*
	(0.21)	(0.31)	(0.34)	(0.21)	(0.22)	(0.23)
ΔX^T_{Act}	0.35	3.29*	3.82*	-0.14	3.33*	4.24*
	(0.57)	(1.31)	(1.48)	(0.56)	(1.39)	(1.65)
$\Delta EmpR_{ct}$	0.35*	0.33*	-	0.34*	0.33*	-
	(0.07)	(0.07)		(0.09)	(0.09)	
Industry	yes	yes	yes	yes	yes	yes
Year	yes	yes	yes	yes	yes	yes
Observations	14,247	14,247	14,247	12,003	12,003	12,003
R^2	0.20			0.99		
Instrument set						
		IV1, IV2 oilTA	IV1, IV2 oitlTA		IV1, IV2 oilTA	IV1, IV2 oitlTA
Sanderson-Windmeijer		Multivariate	F tests			
ΔR_{ct}		35.33	35.1		415.45	451.72
ΔX^T_{Act}		51.23	51.57		21.94	20.47
P-values for S-W tests						
ΔR_{ct}		0.000	0.000		0.000	0.000
ΔX^T_{Act}		0.000	0.000		0.000	0.000

Note: Standard errors, in parentheses, are clustered at the economic region level. The asterisk (*) denotes significance at the 5% level. All models are estimated using 55 economic region by 102 industry cells in four year differences.
ΔR_{ct} is average rent in the economic region. ΔX^T_{Act} is the change in: the probability a person from location c commutes to an ER province times the average rent of commuters in ER provinces. $\Delta EmpR_{ct}$ is the change in the economic region level employment rate.

Table 7: Log Wage and Earnings Levels, Canada and the US

Year	Hourly Wage CPS US	LFS Canada	Weekly Wage Census US	Census Canada	Census US	Annual Earnings Census Canada	Tax Data US	Tax Data Canada
2000	2.69	2.85	6.48	6.56	10.34	10.32	-	10.3
2005	2.68	2.85	6.44	6.6	10.29	10.36	-	10.31
2010	2.68	2.92	6.4	6.65	10.23	10.38	-	10.35

Columns 1 and 2 are based on the same real hourly wage data from the LFS for Canada and CPS for the US used Figure 1. For the remaining columns, the Canadian data is from the 2001 and 2005 Censuses and the 2011 National Household Survey. The US data is from the 2000 US Census and the 2005 and 2010 American Community Surveys. All data is for 20-54 year olds.

Table 8: Estimation Results, Full Country Sample, Including 7x7 Matrices

	OLS	IV		OLS	IV
Variables	(1)	(2)		(3)	(4)
ΔR_{ct}	2.80*	3.21*		2.63*	2.03*
	(0.36)	(0.42)		(0.40)	(0.50)
$\Delta EmpR_{ct}$	0.41*	0.40*		0.42*	0.44*
	(0.10)	(0.10)		(0.10)	(0.09)
$\Delta ERshare_{ct}$	-	-		0.85	7.62*
				(0.99)	(2.40)
Industry	yes	yes		yes	yes
Year	yes	yes		yes	yes
Observations	18,568	18,568		18,568	18,568
R^2	0.50			0.50	
Instrument set		IV1 IV2			IV1 IV2 OILER
Sanderson-Windmeijer		Multivariate		F tests	
ΔR_{ct}	-	626.06		-	134.38
$\Delta ERshare_{ct}$	-	-		-	29.85
P-values for S-W tests					
ΔR_{ct}	-	0.00		-	0.00
$\Delta ERshare_{ct}$	-	-		-	0.00

Note: Standard errors, in parentheses, are clustered at the economic region level. The asterisk (*) denotes significance at the 5% level. All models are estimated using 73 economic regions by 102 industry cells in four year differences.

ΔR_{ct} is average rent in the economic region. $\Delta EmpR_{ct}$ is the change in the economic region level employment rate. $\Delta ERshare_{ct}$ is the change in the component of R_{ct} that corresponds to the ER sector.

Table 9: Decomposition of Resource Boom Wage Effects in Canada

Direct effect within ER provs	0.0024
Add in spillovers within ER provs	0.026
Add in commuting spillovers	0.034
Include labour supply effect	0.036
Include induced demand in non-ER provs	0.037
Actual Change	0.076
Proportion Explained	**0.49**

Table 10: Estimation Results: U.S

	Commute to Oil States				Commute to Housing Boom Cities			
	(1) OLS	(2) OLS	(3) 2SLS	(4) 2SLS	(5) OLS	(6) OLS	(7) 2SLS	(8) 2SLS
ΔR_{ct}	0.75**	0.72**	2.57**	2.59**	0.79**	0.77**	3.10**	2.50**
	(0.081)	(0.081)	(0.65)	(0.65)	(0.11)	(0.11)	(1.18)	(0.97)
ΔX^T_{Act}		5.26**		8.34**		5.10**		13.4**
		(1.42)		(3.80)		(1.68)		(5.69)
$\Delta EmpR_{ct}$	0.19**	0.20**	0.51**	0.48*	0.16**	0.17**	0.17	0.39
	(0.034)	(0.034)	(0.25)	(0.25)	(0.065)	(0.065)	(0.41)	(0.34)
Observations	40529	40529	40529	40529	20058	20058	20058	20058
R^2	0.055	0.056	.	.	0.074	0.074	.	.
Fixed Effects:								
Ind.× Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Instrument set:			IV1-IV3	IV1-IV3			IV1-IV3	IV1-IV3
				$\hat{\psi}_A\times$ Oil$\times$ t				$\hat{\psi}_A\times$ Siaz$\times$ t
First-Stage:								
F-Stat. p-value:								
ΔR_{ct}			0.00	0.00			0.01	0.00
ΔX^T_{Act}			0.00	0.00			0.00	0.00
$\Delta EmpR_{ct}$				0.00				0.00

Notes: Standard errors, in parentheses, are clustered at the county level. (**) and (*) denotes significance at the 5% and 10% level, respectively. All models estimated on a sample of U.S counties using the 2000 Census and the ACS 2006/05, 2010/09, and 2014/15. The dependent variable is the regression adjusted city-industry log wage, and cells with less than 20 observations are excluded. Columns 1-2 and 5-6 are estimated via Least Squares and Columns 3-4 and 7-8 are estimated via Two Stage Least Squares. The bottom panel of the table shows the results of the first-stage statistics for the excluded variables of the 2SLS procedure for columns 3-4 and 7-8.