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

David A. Green[a], Rene Morissette[b], Ben M. Sand[d], Iain Snoddy[d]

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

Abstract

Since 2000, US real average wages stagnated or declined while Canadian wages increased. We investigate the role of the Canadian resource boom in explaining this difference. We focus on wage spillovers to non-resource workers through a bargaining channel. We find that long-distance commuting to resource regions had substantial spillover effects on non-commuters in sending regions. Through spillovers, we account for 49% of the increase in the real mean wage in Canada between 2000 and 2012. We also find long-distance commuting effects in the US. We conclude that long-distance commuting integrates regions, spreading benefits and costs of booms across the economy.

Keywords: Wages, Resource Boom, Inequality

1. Introduction

Since 2000, the Canadian and US labour markets have diverged substantially. Between the early 2000's and 2015, the real hourly wage for US prime age workers fell by 2 to 3 percent for males and was generally flat for females. In contrast, in Canada, the real hourly wage rose by nearly 10% for males and 15% for females over the same period. Similarly, the employment rate for high school or less educated males and females fell substantially over this period in the US but were flat for Canadian men and rose for Canadian women. Inequality has also followed different paths in the two countries, with the log 90-10 wage differential rising by over 10% for both men and women in the US but being essentially unchanged in Canada.

Our goal in this paper is to understand the source of these differences. We argue that a substantial part of the success in the Canadian labour market can be explained by the impact

☆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, Dalhousie University, University of Vienna, and Central European University.

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

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. But to an empirical economist who has worked with shift-share type decompositions, the claim might seem unlikely. In those decompositions, the effect of the boom on the mean wage in the economy reflects employment growth in the high-paying resource sector (the 'between' component) and increases in wages for both old and new workers in that sector (the 'within' component). 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, the between and within components for the resource sector can never be large. This is because the shift-share decomposition focuses on changes for workers in the resource sector itself.

Given this logic, the resource boom can only provide a sizeable explanation for Canadian wage increases if it had spillover effects on the wages of other workers. 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, McKinnish, and Sanders (2005),Michaels (2011), Marchand (2012), Weber (2012), Jacobsen and Parker (2016), Allcott and Keniston (2014), Feyrer, Mansur, and Sacerdote (2017) and Bartik, Currie, Greenstone, and Knittel (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. We complement this literature by investigating a specific spillover mechanism that has the potential to imply large effects that are not only widespread in the local economy but are also transmitted to other, non-resource regions through a migration channel.

Our core idea, following on Beaudry, Green, and Sand (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.

[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).

(2012)'s work 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.

Two points make the resource boom a plausible source of spillovers to the broader Canadian wage structure. First, both Fortin and Lemieux (2015) and Marchand (2015) provide evidence of substantial spillovers within the economies with the highest concentration of employment in the extractive resource sectors – Alberta, Saskatchewan, and Newfoundland. Marchand (2015) compares resource and non-resource communities within the ER provinces and demonstrates spillovers within the local resource communities while Fortin and Lemieux (2015) use provincial level variation, showing spillovers that incorporate those local spillovers and spillovers across communities within the ER provinces. We replicate those results in the second section of the paper. Second, there were links between multiple communities in Canada and the resource intensive areas through migration and long-distance commuting, in which workers maintained a home in a non-resource community and flew back and forth to work in the oil fields for weeks at a time. We argue that the availability of this commuting option increased the bargaining power of workers in all sectors, spreading the effects of the boom both geographically and sectorally to other parts of the economy. Thus, our work complements earlier research both by providing a specific channel through which the spillovers they measure can operate and by looking at impacts beyond the borders of extractive resource intensive regions.

At the heart of our approach is a rich administrative data set that includes the universe of Canadian individual tax filers between 2000 and 2012. These data include tax filers' main tax form (T1) and job specific tax forms (T4's). Importantly, for our purposes, the former contains the address from which an individual files their taxes and the latter records the province of the firm for a particular job an individual worked. Using this information, we can identify 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 a sub-provincial level (Statistics Canada defined Economic Regions), we compare 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 with different levels of 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.

These wage effects could arise 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 address this problem using Bartik-type instruments suggested by an extended version of the models in Beaudry et al. (2012) and Beaudry, Green, and Sand (2014). The instruments 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.

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 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. 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. Thus, in the end, we provide an explanation for a substantial portion of the success in the post-2000 Canadian labour market but no new insights into the decline in the US labour market over this period. We believe it is helpful to provide an explanation for Canadian exceptionalism since comparisons across countries are often used to assess common forces such as trade and technological change. In Canada's case, trade – at least in the form of expanding resource exports – provided wide spread benefits to workers across the country. At the same time, those wage benefits could have implied Dutch Disease type effects on investment and production in other sectors as they faced increased wage costs even in parts of the country that were not receiving direct demand benefits from the resource boom.

The paper proceeds in 8 sections not including the introduction. In the second 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 section 3, we calculate the impact of the boom from a standard shift-share exercise and then provide an heuristic description of a model with spillovers between sectors and across regions. In section 4, we set out our empirical specification and discuss identification issues. Section 5 contains a description of our main Canadian data set, with results from that data found in section 6. In section 7, we present a decomposition of overall Canadian wage movements based on our estimates from section 6. Section 8 contains estimates using US Census data and section 9 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 (left panel) and women (right panel). In the figures that follow, we normalize series to 0 in 1997 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.

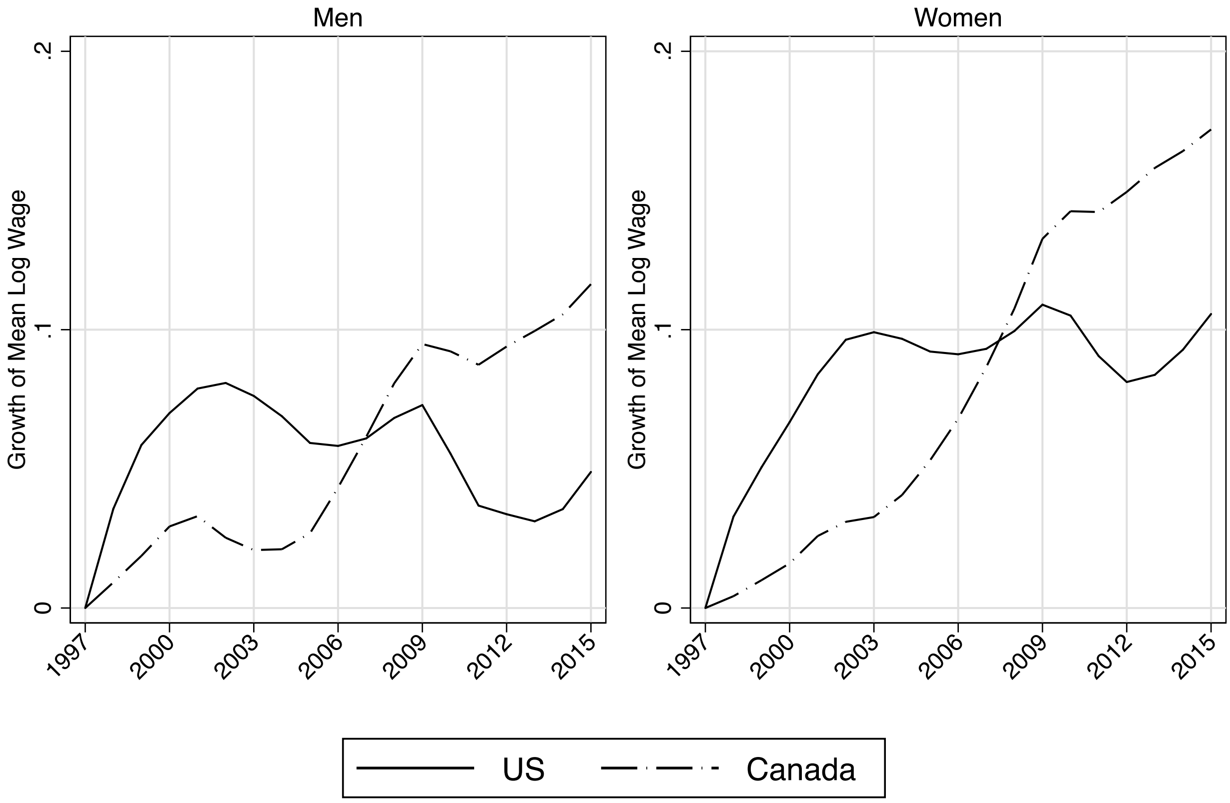

Figure 1: Real Mean Log Hourly Wages, Canada and the US, 1997-2015

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 figures presented in Appendix A, we show that the 90th percentiles show very similar increases in the two countries over our time period while the 10th percentiles follow a similar pattern to the means in figure 1 with declines for the US and growth for Canada after 2003. Thus, after 2003, the whole Canadian wage distribution moves up together while wage gains are concentrated at the top of the distribution in the US.

As discussed in the introduction, the Canadian public and policy makers would likely attribute the relative strength of Canadian wages in this period to the resource boom. This point is also raised in Green and Sand (2015) and Fortin and Lemieux (2015), both of which argue that shifts in Canada's wage structure in the last 20 years have an important regional

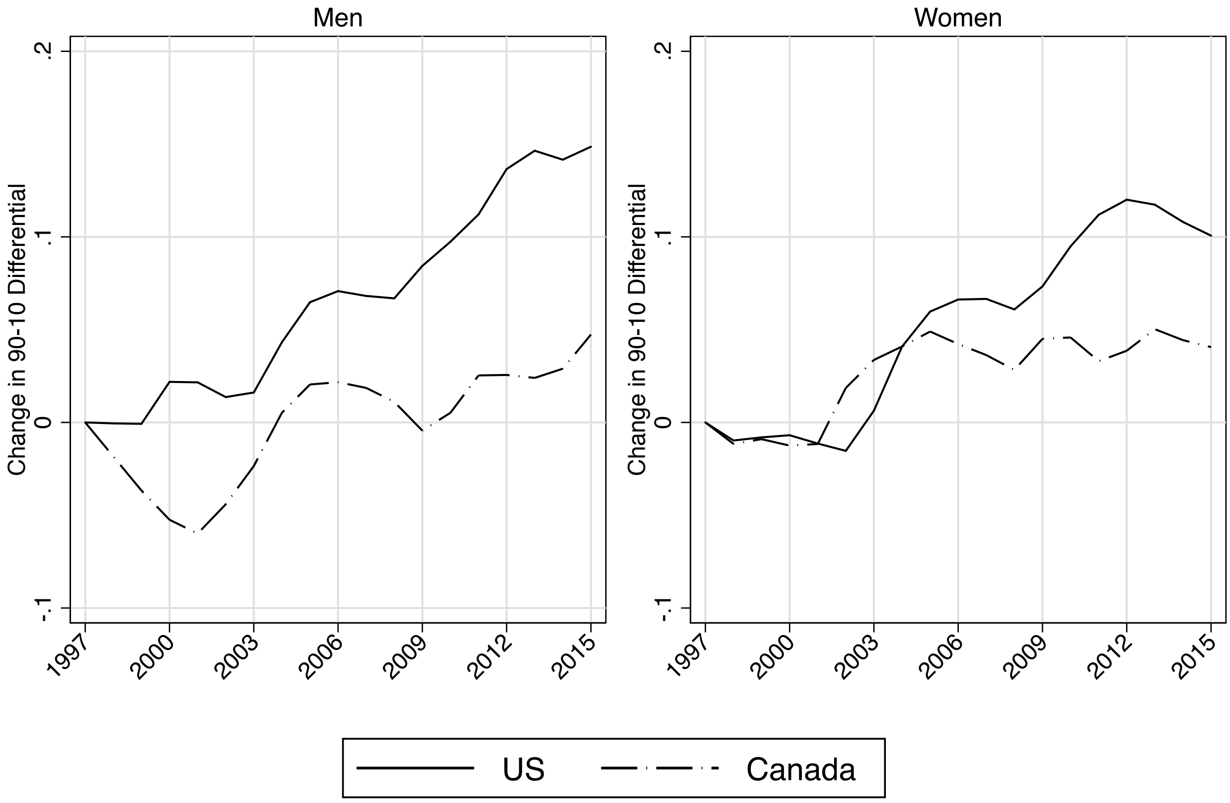

Figure 2: Log 90-10 Hourly Wage Differentials, Canada and the US, 1997-2015

dimension. That is revealed in figure 3 in which we plot real mean log wages for three regions in Canada. The first is what Fortin and Lemieux (2015) call the Extractive Resource (ER) sector: Alberta, Saskatchewan, and Newfoundland.[2] The second is Ontario – which contains Canada's industrial heartland as well as its main financial centre. The third is the Maritime Provinces (PEI, Nova Scotia and New Brunswick), which are perennially poorer provinces located on the Atlantic coast. 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 experience of the Maritime provinces lies between the two and is intriguing in that it shows a wage increase that lags slightly behind that in the ER provinces in spite of not having substantial extractive resources.

[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, Chan, and Lu, 2015).

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 4. 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.[3] 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 constructing our instruments in the next section.[4]

3. Wage Decompositions and Spillovers

The figures in the previous section provide a *prima facia* case that Canadian wage movements are related to the resource boom, but can the boom actually account for a large share of the rise in Canadian mean log wages documented in Figure 1? On average, over our time period, the ER sector made up 1.5% of employment in Canada compared to 0.55% in the US. In Alberta, 8% of workers are in the ER sector compared with 2.2% in Texas.[5] Thus, the resource sector is more salient in Canada. Nonetheless, one might doubt that one sector could cause such large changes in the average wage, even within 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 mean wages that occurred in Alberta over this time period, the apparent 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.

[3]Plotting the oil price over a longer period shows that it is essentially flat between approximately 1983 and 2003 and that the increase in the figure between 1998 and 2000 reflects a recovery from a temporary drop rather than the beginning of the boom. That the boom doesn't start until 2003 is captured in a figure plotting the proportion of employment in extractive resource industries shown in Appendix A. That proportion is actually declining until about 2000 and starts to turn sharply upward around 2003. This corresponds to common definitions of when the most recent boom started.

[4]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). We show figures for different education groups in Appendix A.

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

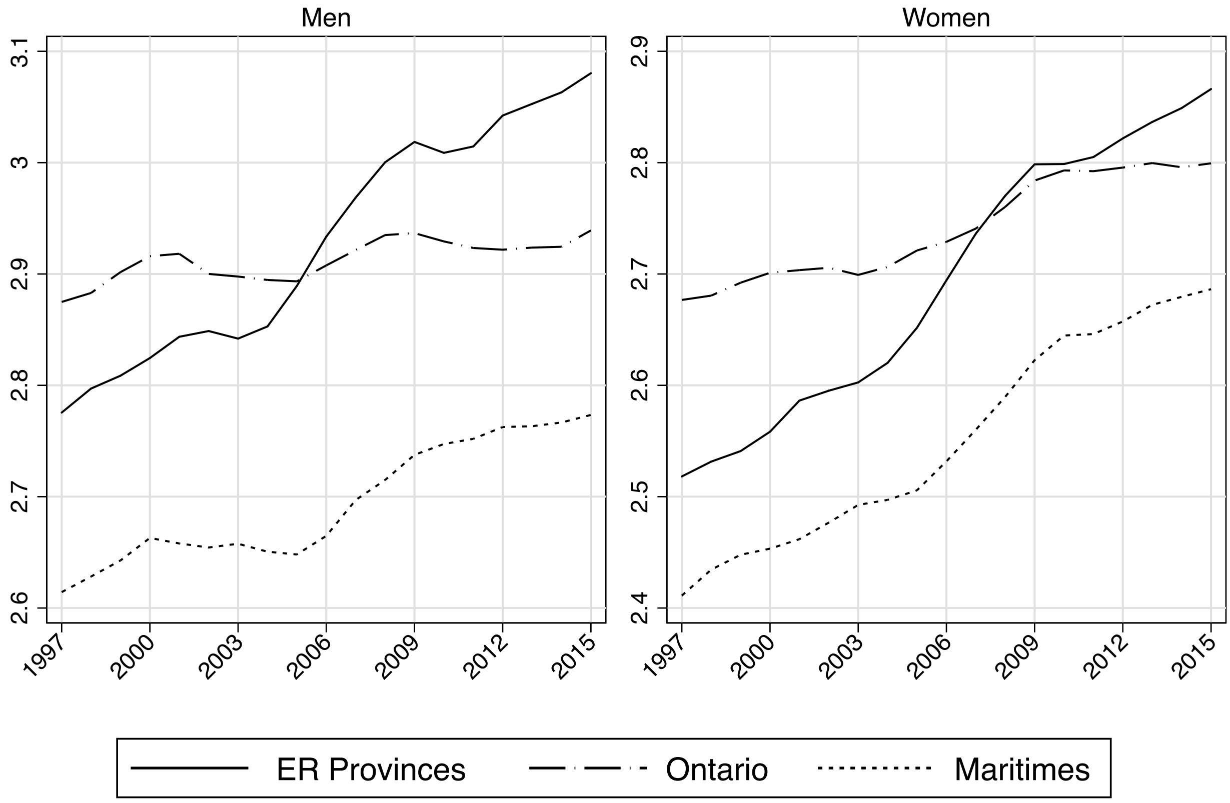

Figure 3: Mean Log Wages by Region, Canada, 1997-2015

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 get new jobs in the sector during the boom. Papers such as Michaels (2011) and Weber (2012) present evidence that discoveries of oil or natural gas in a locality lead to increases in per capita income and employment outside the resource sector (relatively large increases in Michaels (2011)'s investigation of historic oil discoveries in the US South and small ones in Weber (2012)'s shale gas example). Marchand (2012) and Marchand (2015) both use variation in the intensity of extractive resource activity across Census Divisions in Western Canada to identify the spillover effects on earnings. Marchand (2015) shows that the whole distribution of individual earnings shifts up in resource versus non-resource intensive regions, with bigger increases at the top and bottom of the distribution. Since Census Divisions are relatively small areas, these findings correspond only to very local spillovers. Marchand shows that these spillovers spread across sectors, with the largest earnings increases in the construction sector and ER sector itself, and the smallest (though still substantial) increases in retail trade. These papers hypothesize that spillovers can occur through direct demand effects and through local agglomeration externalities, though none of them provide specific evidence on channels.

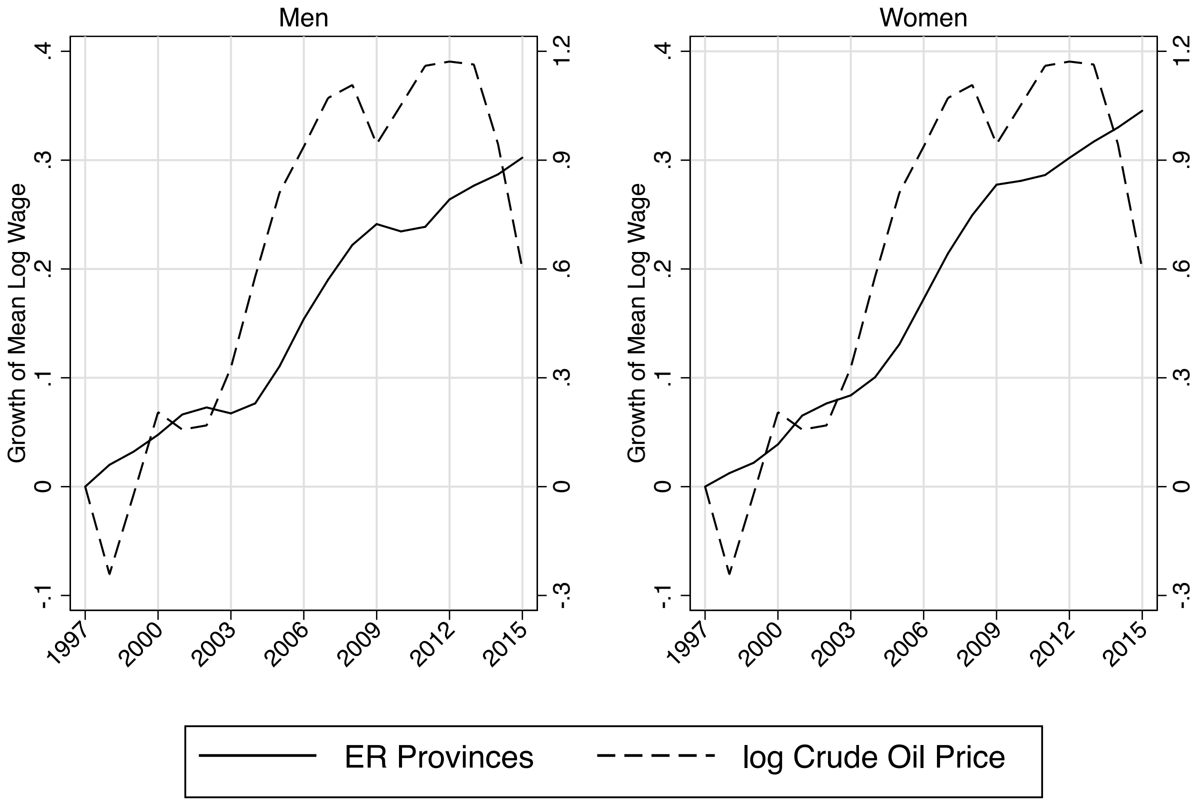

Figure 4: Crude Oil Price and Mean Log Wage in ER Provinces, 1997-2015. Left Scale: Mean Log Wage. Right Scale: Oil Price

Another form of externality is suggested by the search and bargaining model in Beaudry et al. (2012). The model is a standard Mortensen-Pissarides type model but with multiple sectors. In this model, search frictions hinder the immediate reallocation of workers both within and across local labour markets. Due to these frictions, when workers and firms meet, there is a potential surplus of the value of what they produce together relative to the sum of the value of their outside options. If the surplus is positive, workers and firms stay in the match and bargain a wage that serves to divide the surplus. Importantly, the wage moves with the relative values of the options outside the match for the worker and firm. Thus, for example, in tight (high employment rate) local labour markets, it is easier for a worker than a firm to find a new match. This means that the value of the worker's outside option is relatively higher and the wage increases as a result.

The point of the Beaudry et al. (2012) model is that the value of workers' outside options also depends on the composition of employment in the local economy. In economies with more employment in high-rent industries, such as the steel or oil and gas sectors, the value of workers' outside option of continued search is higher since that includes the possibility of obtaining one of the high-rent jobs. Thus, consider two identical workers, both in construction, but one in Alberta with easy access to oil sector jobs and the other in New Brunswick where the main local employer offers low paying call-centre jobs. According to the reasoning of this model, the construction worker in Alberta can bargain a higher wage. Beaudry et al. (2012) show that these implications hold true in US-city level data and that a key over-identifying restriction from the model cannot be rejected. It is important to keep in mind that what matters for bargaining is not simply wage differentials across sectors but rents. If, for example, the higher wages in the oil sector were simply compensating differentials for dangerous work, then workers would not be able to point to those jobs as providing them higher value and therefore could not use them to bargain a higher wage.[6]

In Appendix B, we set out a variant of this model that is extended to allow for the possibility that workers may not only search locally, but may have mobility options. In particular, we consider an environment in which the economy has two geographically separate areas: one region has oil and gas reserves, and we call this the ER province; and the other region, without reserves, is the non-ER province. We envision the latter region as having several local labour markets, indexed by c. Our main focus is on wage determination in cities in the non-ER province during the oil price boom. For this reason, the model we present is partial equilibrium in nature (we take the employment rate and the migration rates as given).[7]

The mobility options available to workers link wage setting across local labour markets, allowing the possibility that resource booms may impact wages in non-resource economies. We assume that unemployed workers, in addition to potentially finding a local job, may receive a mobility shock – the option to either permanently migrate or commute to the ER province. As in standard Rosen-Roback (1979; 1982) geographic equilibrium models, when the option to permanently move arises, workers decide whether or not to migrate by comparing the value of living in the ER province to remaining in c, net of migration costs. If a worker takes this option, we assume that the worker joins the pool of unemployed workers in the ER province. The second mobility option is for a worker to take up long distance commuting. With this option, a worker continues to live, pay housing costs, and enjoy the local amenities in c. However, the benefits of commuting depend on the location specific costs of commuting. These costs vary across c due to the distance of c to the ER

[6] 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 rent variable that we describe in the next section should not affect wage setting within sectors.

[7] We define the relevant local labour markets as Statistics Canada's Economic Regions and treat migration in a relatively reduced form way. In contrast, Manning and Petrongolo (2017) use UK vacancy, unemployment, and mobility data at a fine local level to investigate the effective size of local labour markets.They find that costs of commuting rise sharply with distance. Established channels of movement such as those to the resource rich regions in Canada could serve to reduce those costs. Certainly, as detailed below, we find substantial evidence of well-used channels. We view our work as complementary to theirs in that we focus more on the wage setting mechanism. In their work, wage spillovers to other regions from a boom in one region happen through reductions in local supply of workers. As we discuss in our empirical specification and results sections, our approach allows for that channel as well as a bargaining channel for spillovers. We find that the bargaining channel is substantially more important.

province or to non-pecuniary factors such as language. As we make clear below, these two mobility options are quite distinct, and can have very different implications for local wage determination.[8]

When an unemployed worker meets a local employer, a bilateral monopoly is formed and workers and firms bargain a wage that is set according to a standard Nash bargaining rule that divides the surplus of the match. Importantly, the bargained wage will depend on workers' value of continued search. In our setting, this value will depend on the local industrial composition, and on the value of the mobility options which will improve with the onset of the oil boom. Because the mobility options improve the value of continued search for all workers, those in non-ER regions who neither migrate nor commute can use these improved options to bargain better wages; that is, the options have a nonrivalrous quality. This can imply spillovers of substantial size across the local economy.

Finally, it is useful to discuss the two mobility options in more detail. If the option to permanently move occurs frequently enough (the mobility friction is not too high), a steady state spatial equilibrium will imply that the value of permanent moves are driven to zero. As in standard spatial equilibrium models, we assume that housing is not perfectly elastically supplied. Thus, an increase in oil prices will initially make the ER province more attractive and they will receive an inflow of workers from non-ER regions. This will simultaneously increase the cost of housing in ER provinces while lowering it in the non-ER regions, and this process will continue until expected utility is equated across the two regions. Thus, in a steady state spatial equilibrium, the option to permanently move plays no role in local wage determination. This is because the permanent move option, provided mobility is high enough, is directed – workers choose to search locally or move and search in the ER province, and workers are indifferent between these options in equilibrium. Commuting, in contrast, is not directed; this option acts like an extra set of industry options for the local market. In addition, workers who commute face the same housing prices and, thus, the mechanisms which equate utility across locations in a spatial steady state concept are absent.[9] Therefore, provided commuting is viable (commuting costs are not too high), the commuting option will always play a role in local wage determination.[10] However, if permanent move mobility

[8]In addition to these options, an individual could also either move permanently or engage in long-distance commuting to a different location in non-ER provinces (e.g., Toronto). We have estimated specifications including the value of these options but the estimated effects are always poorly defined. This likely arises because nearly all of the non-ER commuters are in the Ottawa region where the local economy straddles a provincial border.

[9]The long distance commuters also are able to maintain family ties. Kennan and Walker (2011) and Zabek (2018) argue that regional migration patterns are strongly influenced by a home bias.

[10]The Canadian unemployment insurance (called EI) could also affect the relative value of the outside options. The hours of work needed to qualify and weeks of benefit eligibility vary with the local unemployment rate in the Canadian system and the worker collects according to the unemployment rate where he resides not where he works. In Appendix B, we discuss the implications of this system for our specification and results. Given that we control for changes in the employment rate in a differenced specification at the local level, any EI effects are controlled for and do not alter our results or interpretation. We also derive a version of the model that includes a benefit that is proportional to the past wage (as is the case in the Canadian system). We show that this magnifies spillover effects since having access to better jobs improves

frictions are large, the value of this option will not be driven to zero and workers can use this option to bargain higher wages locally. In our baseline work, we assume worker flows are sufficient to equate expected utility across markets, but we also empirically examine the relevance the the permanent move option below.

4. Empirical Specification

We are interested in the potential impacts of the resource boom on the workers in non-ER provinces who do not migrate or commute to ER provinces – the long-distance spillovers. For that reason, we use the model from Appendix B to derive an empirical specification for wages in industry i in a non-ER province city c:

$$\Delta \ln w_{ict} = \beta_{0it} + \beta_1 \Delta R_{ct} + \beta_2 \Delta X_{ct}^T + \beta_3 \Delta Emp_{ct} + \xi_{ict}, \tag{1}$$

where $\Delta x_t = x_t - x_{t-1}$. This equation relates local, within-industry wage movements to two key variables in the model that capture, in part, the strength of workers' outside options in location c. The first of these, R_{ct}, is an index that captures the extent at which there are 'good jobs' available locally. This index, which we refer to as average local rent, is constructed as $R_{ct} = \sum_j \eta_{jct} \cdot \nu_{jt}$, where the weights, η_{jct}, are local employment shares and ν_{jt} are national-level wage premia.[11] A high value of this index indicates that local employment is concentrated in high-paying sectors.

The second variable, X_{ct}^T, captures the value of the commuting option for workers originating in c. This variable is calculated as $X_{ct}^T = q_{ct}^T \cdot \sum_j \eta_{jct}^T \cdot \nu_{jt}$, which, again, is function of an index, $\sum_j \eta_{jct}^T \cdot \nu_{jt}$, that reflects average rents available in the ER-provinces for commuters from c, where η_{jct}^T are the industrial proportions (in their ER jobs) of c commuters. This index is scaled by q_{ct}^T, which captures the probability that a worker has the option to commute. This probability is, in part, a function of the local commuting costs, with lower commuting cost localities having a larger value of q_{ct}^T, all else equal.

Before discussing the interpretation of the impact of each of these variables, it is useful to first comment on two other aspects of the wage equation. First, the β_{0it} term in (1) has i and t subscripts and is captured in our estimation with a complete set of year×industry fixed-effects. Second, the specification includes the local employment rate, Emp_{ct}, as a control. These two elements reflect the fact that, in the model, wages are determined by productivity and by factors that influence a worker's bargaining position. The year×industry effects capture productivity factors, most notably the national level prices for industrial output. The employment rate captures labour market tightness, with the relative bargaining position of workers increasing in that tightness. Thus, its coefficient, β_3 is predicted to have a positive

a worker's outside option both because it is good in itself and because better paying jobs lead to higher benefits when the jobs end. The commuting option is particularly enhanced by this system since workers can get benefits based on the high wages in Alberta over the extended weeks of qualification relevant in the Maritimes.

[11] In our random search model, we show that workers meet different sectors in proportion to the size of the sector in the local economy.

sign. Similarly, β_1 and β_3 are predicted to be positive since they capture the effects of increased values for outside options.

Given this specification, the estimated effects of our key variables exploit within-industry, over-time variation, conditional on the local employment rate. Intuitively, this means that we identify the impact of local average rent, β_1, by comparing wage changes for workers in the same industry in two different cities that are experiencing different changes in their industrial composition. The fact that we control for the employment rate implies that the estimated effect stems purely from compositional shifts, and not either from the level of labour demand or shifts in local labour supply. Our interpretation of β_1 is that this coefficient captures a spillover effects via wage bargaining, due to an improvement in workers' outside options.[12] As we explain in Appendix B, our model implies that wages in a local industry are related to average wages for the city as a whole in the form of a standard reflection problem. If a locality's industrial structure changes there will be a direct effect on average wages, representing the mechanical relocation of workers across sectors. This will be accompanied by indirect effects as the shift in average local wages impacts bargaining in all sectors, initiating a cycle of wage adjustments inherent in the reflection problem. The mechanical effect is a between-industry effect, while the indirect or spillover effects are within-industry effects. Importantly, our specification focuses on within-industry wage movements and uses national-level premia to solve reflection issues; as a result, we show that the coefficient β_1 captures the extent of spillovers from local industrial composition shifts. For example, we identify the bargaining effect by comparing wage changes for construction workers in Hamilton, with the loss of its high wage rent steel sector, to construction workers in, say, Moncton, without the loss of such a sector. The core idea is that, before the loss of the steel sector, construction workers in Hamilton will be able to bargain higher wages than those in Moncton because their outside option included the possibility of getting high wage steel jobs, but this advantage disappears along with the steel sector.

Similarly, β_2 captures the impact on local wages through bargaining from shifts in the value of commuting to ER provinces. The value of the commuting option depends, in part, on national-level wage premia and the industrial composition in the ER provinces, which are common across non-ER localities. Variation in the value of X_{ct}^T across non-ER cities will partly be driven by location specific commuting costs, which, for example, may be a function of distance. Identification of β_2 comes from comparing the wages of workers in the same industry across cities that experience different changes in X_{ct}^T when the ER provinces' industrial composition improves. It is worth emphasising that (1) focuses on the wages of workers not directly involved in commuting (ie, those who actually work in c). Thus, by its very nature, the effect captured by β_2 is an indirect or spillover effect from the resource provinces to non-resource communities, linked via commuting. As we discuss in more detail below, a key to our identification strategy will be to exploit pre-oil boom

[12] Our interpretation relies on the idea that ν_{jt} represent rents; that is, wage differences across industries for workers of the same skill, and not productivity differences between workers 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. In the model, wage differences between industries arise because of differential prices of industrial goods.

commuting intensity, which we interpret as being a function of location specific commuting costs.

Finally, equation (1) only includes the value of the commuting mobility option, as given by X^T_{ct}. As discussed in section 3, workers also have the option to permanently move from non-ER locations to the ER provinces. The value of such an option will depend on whether or not a spatial equilibrium has been reached. Equation (1) captures a baseline scenario of a steady state equilibrium, and thus the value of permanent moves are zero and therefore not included. If a steady state spatial equilibrium is not reached, however, then equation (1) will contain another term capturing the value of the permanent move mobility option, which we denote X^P_{ct}.[13] In our empirical section below, we assess the relevance of the permanent move option by estimating a version of equation (1) that includes this term.

With the wage specification at hand, 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, an increase in the price of oil will show up as an increase in the size of the resource extraction sector and/or an associated increase in the wage premium associated with that sector, improving the average local rent in ER provinces. This, in turn, generates increases in bargained wages in other sectors in ER provinces through the mechanisms outlined above. 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 locations that are farther away.

The first effect of the increase of wages in the ER provinces will be to induce some workers in non-ER provinces to migrate or commute to jobs in the ER provinces. This results in a reduction in the supply of workers in c, inducing an increase in wages there. In the model, this immediate impact arises because the employment rate, Emp_c, rises and tighter labour markets benefit workers in wage bargaining. This first, labour-supply channel (whether in the short or long run) is not what is being captured in the β_1 and β_2 coefficients in (1), because our specification includes a control for the change in the employment rate. 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 the ER provinces, 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.

Additionally, the increase in wages in the ER provinces will increase the value of the commuting option for those bargaining with employers locally in c, increasing local wages as described above. Workers will also be induced to migrate permanently and, as long as mobility frictions are not too high, workers will pursue this option until is value is driven

[13] We construct X^P_{ct} in a similar fashion to X^T_{ct}. In particular, $X^P_{ct} = q^P_{ct} \cdot \sum_j \eta^P_{jct} \cdot \nu_{jt}$, where q^P_{ct} captures the probability that a worker from c will have the option to permanently move to an ER province, and η^P_{ct} captures the industrial composition of those who permanently move.

to zero and a spatial equilibrium holds. While we maintain this assumption in our baseline work, in Appendix B we show that if a spatial equilibrium fails to hold, we can add a term to equation (1), comparable to X_{ct}^T, that represents the value of this mobility option. This allows us to empirically assess whether a spatial equilibrium holds as this variable should have no impact on wage determination if workers are indifferent between locations.

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 resulting shift in industrial composition in c 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 wage effects. However, as part of our decomposition exercise in section 6, we will use input-output tables to assess the size of the latter channel.

4.1. Endogeneity and Identification

Another key aspect of the wage specification, as we outline in the theory appendix, is that the error term is a function of local and ER province cost shocks. This helps to highlight the sources of potential endogeneity and inform our identification strategy. Specifically, as we outline in Appendix B, the error term is a function of changes in local industry cost shocks and, as as discussed in Beaudry et al. (2012), these shocks, in part, determine the local employment shares that form one of our key regressors, ΔR_{ct}, the change in local average rent. Additionally, the error term contains changes in cost shocks in the ER provinces, which impact X_{ct}^T through both the commuting intensity and commuters' industrial employment. In this section, we outline an instrumental variables approach in an attempt to overcome these endogeneity issues.

In order to form instruments for local average rent, we adopt an instrumental variables approach that is based on that used in Beaudry et al. (2012). In particular, one can decompose ΔR_{ct} as:

$$\Delta R_{ct} = \sum_j \left(\eta_{jct} - \eta_{jct-1}\right) \cdot \nu_{jt-1} + \sum_j \eta_{jct} \cdot \left(\nu_{jt} - \nu_{jt-1}\right). \tag{2}$$

This decomposition says that the change in local average rent can be broken down into changes that arise because of changes in the industrial composition, holding the industrial wage premia constant (a between-industry component) and changes that arise because of changes in the premia, holding the local industrial composition constant (a within-industry component). Beaudry et al. (2012) argue that each of these components can be seen as the basis for a valid instrument. In particular, they construct Bartik-style instruments based on each component by replacing the period t industrial share with a prediction based on the national-growth rate of employment in each industry and period $t-1$ local industrial composition.[14] This gives the following instruments:

$$IV1_{ct} = \sum_j (\hat{\eta}_{jct} - \eta_{jct-1}) \cdot \nu_{jt-1} \quad \text{and} \quad IV2_{ct} = \sum_j \hat{\eta}_{jct} \cdot (\nu_{jt} - \nu_{jt-1}) \tag{3}$$

Beaudry et al. (2012) provide a detailed discussion of the conditions under which these instruments are valid. Here, we 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 variation. That means that what we are concerned with is the cross-city correlation between $IV1_{ct}$ or $IV2_{ct}$ and ξ_{ict}. The cross-city variation in $IV1_{ct}$ and $IV2_{ct}$ comes from cross-city differences in the η_{ict-1}'s.[15] Thus, for our instruments to be valid, we must assume that cross-city differences in start-of-period industrial composition are uncorrelated with the error term – a standard condition for Bartik-style instruments (see Goldsmith-Pinkham, Sorkin, and Swift (2018)). In our model, this requires start-of-period industrial composition not to predict general productivity growth in a city. This may not seem credible – at least *ex post*, it seems possible to identify types of industrial mixes that predict general growth in an area.

In response to this, Beaudry et al. (2012) suggest using both instruments and performing an over-identification test. In particular, each instrument exploits different variation in ΔR_{ct}, either that stemming from the within component ($IV1_{ct}$) or the between component ($IV2_{ct}$), but the model implies that each source of variation should have the same impact on wages. 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, and there is a correlation between the η_{ict-1}'s and the error term, the two instruments will 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, which forms the bases of the over-identification tests we perform below.[16]

Next, we form an instrument for ΔX^T_{ct} that exploits changes in oil prices over the boom, which we believe is the main driver of changes in ER province wages over the period we examine. This helps focus our discussion on the benefits of resource booms for non-resource

[14] In particular, we construct a predicted level of employment, $\hat{E}_{ict} = E_{ict-1} \cdot (1 + g_{it})$, 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, $\hat{\eta}_{ict}$

[15] To see this, note that $\hat{\eta}_{ict}$ gets its cross-city variation from η_{ict-1}.

[16] In $IV1_{ct}$ the potentially offending correlation is weighted by national industrial growth rates and the start of period industrial wage premia, and in $IV2_{ct}$ 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. In an appendix, Goldsmith-Pinkham et al. (2018) argue that our over-identification test is not consistent. In a note, available upon request, we show that their argument is true only in a knife-edge sense and our over-identification test is useful.

economies by explicitly focusing the source of the boom. In addition, as Kline (2008)'s argues that workers respond more to the probability of getting a job as signalled by oil price movements than to wage changes. In order to obtain cross-city variation in our instrument, we combine the oil price changes with pre-oil boom commuting intensity. As discussed above, we interpret pre-oil boom commuting intensity to be, in part, driven by differences in the cost of commuting from each c to the ER provinces. Thus, our instrument is constructed as:

$$IV_{ct}^{X} = \hat{q}_{c2000}^{T} \cdot \Delta \ln p_t^{oil}.$$

The identifying variation used in IV_{ct}^{X} can be thought of in terms of a difference-in-difference framework; oil price changes act as a continuous treatment and start-of-period commuting intensity determines exposure. Our identification assumption is that start-of-period commuting costs are uncorrelated to changes in local and ER province costs shocks that make up the error term. Intuitively, this assumption states that communities with low-commuting intensity would have experienced the same wage changes as those with high-commuting intensity, in the absence of the oil price boom (conditional on our other controls).[17]

5. 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 individual tax form (the T1) for all individuals in Canada.[18] 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.[19] 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.

[17] To address the endogeneity of ΔEmp_{ct}, we also tried 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, $IV3_{ct} = \sum_j \eta_{jct-1} g_{jt}$. We discuss the use of this instrument in our results section.

[18] The T1PMF does not include the 5% of taxfilers who file late but otherwise contains the universe of tax filers. Messacar (2017) investigates late filers and finds that including or removing them does not alter the observed earnings distribution.

[19] Economic Regions around CMAs are approximately the same as commuting zones but in rural areas, they are agglomerations of commuting zones. We chose the larger areas in order to allow for detail in the set of industries. Rural commuting zones often have samples that are too small to allow anything but very crude industry groupings.

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 reported location refers to the place of work not the location of the firm's head office. Since 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. 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).

5.1. Implementation

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.[20] 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.

In order to empirically implement our framework, we must estimate the national-level industrial premia to construct our rent variables. In our theoretical set-up, workers are

[20]We choose these periods in order to divide our data into even medium-term sub-periods. They also roughly correspond to pre-boom, rapid boom, and mature boom periods.

homogeneous and wage differences across industries are rents. To account for worker heterogeneity while focusing on the within-skill implications of our framework, we use regression adjusted industrial wage premia. To construct these, we run regressions, separately by year, 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. The coefficients on the industrial dummies correspond to the ν_{it}'s, and are used along with observed local industrial employment distributions to construct R_{ct} and X_{ct}^T.

Even after these adjustments, we have some concerns that the ν_{it}'s could still reflect differences in productive characteristics such as education (which is not available in our administrative data) rather than rents. To address this, we construct proxies for worker skill that exploit the panel nature of our administrative data. In particular, we include in our premia-adjustment regressions dummies representing cells in 7×7 matrices defined by the 5 earnings quintiles, self-employment status, and not-employed status in year $t-1$ on one dimension and the same 7 categories in year $t-2$ on the other dimension. The earnings quintiles in these matrices are gender×age×province-specific, i.e. are constructed for 180 groups defined by province of residence, gender, and 9 age groups (20-24, 25-29, 30-34, . . . ,60-64). This means that current year industry effects are 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 the matrix cell indicators.

Similarly, to construct our dependent variable in (1), we also adjust city-industry wages via regression. In particular, we estimate the same specifications as described above, but include a full set of city-by-industry dummies. The coefficients on these dummies are the mean log earnings for an i-c cell purged of individual characteristic effects, and the difference in these coefficients over time form our dependent variable, $\Delta \ln w_{ict}$. As before, we do this with and without the 7×7 matrices representing earnings groups for the preceding two years.

It should be emphasized that when constructing our dependent variable, we use data only on workers who are not commuters or migrants. That is, we use data only on workers who actually work in c in time t. One potential issue is that the composition of this group might change over time as migration or commuting intensity increases due to the oil boom. Thus, we also estimate our main specifications with an alternative sample that addresses any issue with changing composition over time. In particular, we estimate specifications that focus on a sample of workers that are residents of c who never commute or migrate during the 2000-2012 period. This addresses a specific selectivity issue – the possibility that observed resident average wages might change because the best workers have moved or started commuting to oil-producing provinces.[21]

[21]Our panel data sample consists of individuals who are aged 22 to 42 in 2000 (34 to 54 in 2012) who are observed in c in all years. This group of residents is, of course, a select group in their own right. Our assumption is that wages can be written as a time varying price per efficiency unit times a person-specific

5.2. Variable variation

In order to understand the source of our identifying variation, it is helpful to see 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 those who live and work in non-ER provinces (Stayers), those who live in Non-ER provinces but work in ER provinces (Commuters), and those who migrate from non-ER to ER provinces from $t-1$ to t (Migrants) in both 2000 and 2012. The last two columns of Table 1 refer to the industrial composition of those in ER provinces.

A few key points stand out from the table. 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, Hurst, and Notowidigdo (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.[22] 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.[23]

Table 1 shows 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

(and constant) number of efficiency units. Thus, wage movements for a consistent set of workers will reflect movements in the price of labour, which is our focus. Since we use 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.

[22] Tsvetkova and Partridge (2016) show that construction employment is particularly strongly related to local changes in oil and gas extraction in the US.

[23] While not shown in Table 1, we are also able to identify commuters and migrants from one economic region in the non-ER provinces to another. The within non-ER provinces commuting 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.

such commuters was 0.1% in 2000 and 0.2% in 2012. More generally, in 2012, there were a considerable number of regions with few commuters (with the median of the distribution of the percentage of workers who commute being 0.52%) but also a significant number with more substantial commuting (with the 75th percentile of the distribution being 1.3% and the 90th percentile being 2.9%). For young men, the 75th percentile is 2.9% and the 90th percentile is 6.5%. The high commuting regions are also spread across the country, with the top ten in terms of proportion commuting including regions in PEI, Nova Scotia, New Brunswick (i.e., the east coast), Manitoba (in the centre), and British Columbia (the west coast). Ontario also has some relatively high proportion areas in the more rural northwestern part of the province. In comparison to commuting, 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. According to the model, 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 are reasons to question the likely impact of permanent migration on bargained wages of the residents.

Table 2 shows the changes in the industry premia (ν_i's) between 2000 and 2012 computed with (first column) and without (second column) the 7×7 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 7×7 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 7×7 matrix.

6. Results

In Table 3, 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 equation (1), derived from theory outlined in Appendix B. The dependent variable is changes in the mean log wage in a (non-ER province) location×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_{ct} – 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_{ct} – 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), and the change in the employment rate in c (ΔEmp_{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 effect of the value of the commuting option is positive (as predicted) and significant at the 5% level, while the estimated effect of the permanent move option value is smaller than its standard error.

In the second column of Table 3, 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 bottom panel of the table presents the p-values from the Sanderson-Windmeijer F-tests, which indicate that we do not face weak instrument problems.[24] The estimated coefficient corresponding to the permanent move to ER provinces option continues to be 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 2SLS estimates omitting 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 7×7 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_{ct}^{T} and ΔR_{ct} but not ΔEmp_{ct}. Following Stock and Watson (2014), we then interpret ΔEmp_{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

[24] For the permanent move option, X_{ct}^{P}, we construct instruments analogous to $IV1$ and $IV2$, scaled by $\hat{q}_{c2000}^{p}$, the proportion of workers who move from c to the ER provinces in the two years prior to 2000. We have experimented with other formulations of instruments for the permanent move option, but all yield similar results.

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 $B \times (1 + \beta_1)$, 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_{ct}) 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. What is important for us is 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.[25] Similarly, when we control for the employment rate, the coefficient on ΔX^T_{ct} reflects bargaining effects and does not include effects resulting from reductions in local labour supply because of the increase in commuting.

[25] 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. IV1 corresponds to changes in rent stemming from changes in industrial composition and IV2 captures changes in rent from changes in industrial premia. The two variables use very different variation but should imply the same effects since, for bargaining, it doesn't matter whether rents decline because high paying industries leave town or because they stop paying a premium. 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.

In column 5, we present a specification in which we drop the employment rate variable. This means that the estimated coefficient on ΔX^T_{ct} 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 4 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 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 3 and the observed value of ΔX^T_{ct} 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 3 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. 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 3 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 5, we present two further sets of estimates. In the first, we do not use the 7×7 matrices in the estimations in which we obtain the industry rents and our dependent variable. We view the 7×7 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.[26] The results including the 7×7 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 5 provide estimates from our main specification (without the permanent migration options and the commuting option to non-ER provinces) estimates by OLS and by 2SLS with and without the employment rate control. These results show the same broad patterns as those in Table 3. In particular, the location rent variable and the commuting to ER provinces effects are positive and generally statistically significant in the 2SLS 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_{ct}^{T}. The estimated effects for ΔR_{ct} are of similar size to those in Table 3 but the coefficients on ΔX_{ct}^{T} are considerably smaller. However, the values for ΔX_{ct}^{T} are larger when the industry

[26]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.

premia are obtained when not controlling for prior earnings history. For all non-ER regions combined, the increase in ΔX^T_{ct} 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 7×7 past earnings matrix. Thus, the results from the two approaches are similar in size.

Our 2SLS estimates to this point attempt to address endogeneity in the sense that higher out-migration regions might be ones suffering economic problems but do not address the potential 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, lower ability workers then we would expect an increase in the mean observed wage with increases in the value of ΔX^T_{ct} simply because of the composition change of the non-commuting workers.[27] 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 5, where we again use our main, more parsimonious specification and use the 7×7 matrix of past earnings. The estimated effect for ΔR_{ct} and ΔX^T_{ct} are again statistically significant and positive but are smaller in size than the matching estimates not using the consistent group presented in Table 3. 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 7×7 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.

7. 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 6, 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).

[27] Recall that ΔX^T_{ct} is partly a function of the proportion of people from the location who commute.

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.[28] 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 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 consistency 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 consistency 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

[28] In all cases, we restrict our attention to 20 to 54 year olds.

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 7. 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 ER provinces 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 in those provinces between 2000 and 2012.[29] 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 16% 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

[29] 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 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 through this induced demand channel 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 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 8. 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.

8. 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 (See Appendix A). The US geographical unit that most closely resembles Economic Regions in the Canadian data is the commuting zone

(Autor, Dorn, and Hanson, 2013). Thus, we use US commuting zones as our 'cities'.[30] We then run our parsimonious specification (the one that does not include permanent moves) with the units of observation being city by industry cells in the non-ER states. As in our Canadian exercise, the dependent variable is the change in the average log weekly wage in the given city by industry cell where that average is obtained from a regression in which we flexibly control for education, gender, experience, race, and immigrant status. We construct ΔR_{ct}, ΔX_{ct}^T and the change in the city employment rate along with the instruments for ΔR_{ct} and ΔX_{ct}^T in the same way as in the Canadian data.[31] More details on data construction are provided in Appendix C.

In the first four columns of Table 9, we present results from our US estimation both with and without the commuting option variable and estimated either by OLS or 2SLS. 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 commuting zone level. Our preferred specification is the 2SLS 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 1.80. This is smaller in magnitude to what Beaudry et al. (2012) find using US Census data over a longer period. However, in column 3, which excludes the commuting option, we obtain results that are very similar to that paper. The coefficient on the commuting option is quite similar to what we observed in Table 3 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.[32]

[30] Results obtained when using a PUMA based measure of geographical areas gave similar results. These results are available upon request.

[31] Thus, X_{ct}^T 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.

[32] To check to make sure our US-Canada comparisons were not confounded by using very different datasets, we re-ran our specifications using Canadian Census data in order to match the US data. We again obtain positive and statistically significant effects of ΔR_{ct} and ΔX_{ct}^T on local wages. The ΔX_{ct}^T is scaled differently from the tax data because the commuting probabilities obtained from the Census data are lower than those in tax data. However, when we combine the mean ΔX_{ct}^T with its associated estimated coefficient, the implied increase is 0.54%, which is similar to the 0.9% obtained from tax data. The results from this exercise are presented in Appendix D.

9. 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 province but get their main earnings from a firm in an extractive resource province. 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, in estimation with US Census data, we find very similar size elasticities (relative to our Canadian estimates) of the option of commuting to extractive resource states on wages in non ER states. 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 one-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 in Cape Breton 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. Manning and Petrongolo (2017) make a similar point in their study of the extent of local labour markets and the way they interconnect.

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, Dorn, and Hanson, 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 wage structures in Canada and the US parted ways after 2000 both because Canadian wages rose and US wages fell. Our results indicate that an important part of the Canadian rise can ultimately be attributed to the resource boom. Since the resource sector is much less salient in the US economy, the explanation for declining wages there needs to be found elsewhere. Technological change has often been advanced as an explanation for weakening labour market outcomes in the US. Investigations of the impact of technological change sometimes take the form of cross-country comparisons. Our results suggest that Canada is not a useful data point in these comparisons since any technological change effects appear to be swamped by resource sector effects after 2000.

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

Autor, D. H., Dorn, D., Hanson, G. H., oct 2013. The China Syndrome: Local Labor Market Effects of Import Competition in the United States. American Economic Review 103 (6), 2121–2168. URL http://pubs.aeaweb.org/doi/abs/10.1257/aer.103.6.2121

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

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.

Goldsmith-Pinkham, P., Sorkin, I., Swift, H., March 2018. Bartik instruments: What, when, why, and how. Working Paper 24408, National Bureau of Economic Research. URL http://www.nber.org/papers/w24408

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

Kennan, J., Walker, J. R., 2011. The Effect of Expected Income on Individual Migration Decisions. Econometrica 79 (1), 211–251.

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

Manning, A., Petrongolo, B., 2017. How Local Are Labor Markets? Evidence from a Spatial Job Search Model. American Economic Review 107 (10), 2877–2907.

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.

Messacar, D., September 2017. Big Tax Data and Economic Analysis: Effects of Personal Income Tax

Reassessments and Delayed Tax Filing. Canadian Public Policy 43 (3), 261–283.
URL https://ideas.repec.org/a/cpp/issued/v43y2017i3p261-283.html

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

Roback, J., 1982. Wages, Rents, and the Quality of Life. The Journal of Political Economy 90 (6), 1257–1278.
URL http://www.jstor.org/stable/10.2307/1830947

Rosen, S., 1979. Wages-based Indexes of Urban Quality of Life. In: Mieszkowski, P., Straszheim, M. (Eds.), Current Issues in Urban Economics. John Hopkins Univ. Press, Baltimore.

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

Stock, J. H., Watson, M. W., 2014. Introduction to Econometrics. Pearson, Boston.

Tsvetkova, A., Partridge, M. D., 2016. Economics of modern energy boomtowns: Do oil and gas shocks differ from shocks in the rest of the economy? Energy Economics 59, 81–95.

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

Zabek, M., 2018. Local Ties in Spatial Equilibrium.

Table 1: Industrial Composition for Residents, Commuters, and Migrants: 2000 and 2012

	Non-ER Provinces						ER Provinces	
	2000			2012			2000	2012
	Stayers	Commuters	Migrants	Stayers	Commuters	Migrants		
Agriculture	1.7	4.1	2.5	1.4	1.9	1.2	2.2	1.4
Mining and Oil	0.6	8.5	4.1	0.6	16.3	6.6	5.0	6.5
Utilities	0.9	0.4	0.4	0.8	0.7	0.7	0.9	1.2
Construction	4.6	18.2	9.3	6.0	33.7	12.1	6.7	9.3
Manufacturing	18.5	5.9	10.6	12.0	2.9	6.9	9.2	6.8
Wholesale Trade	5.4	3.4	4.6	5.0	2.9	3.9	5.3	4.7
Retail Trade	9.8	5.0	9.9	10.4	2.6	9.7	10.4	9.8
Transportation	4.5	7.7	5.7	4.4	7.2	5.2	5.2	4.8
Culture	2.5	1.3	1.9	2.3	0.5	1.4	2.4	1.6
Finance	4.9	1.9	3.3	5.1	0.8	2.5	3.6	3.2
Real Estate	1.3	1.5	1.4	1.6	2.0	1.9	1.6	1.8
Professional	4.9	6.2	5.5	5.9	6.6	7.1	5.3	6.6
Management	0.5	0.9	0.5	0.8	0.4	0.6	0.6	0.6
Waste	4.0	5.1	5.4	4.9	5.9	6.8	3.4	3.9
Educational	7.9	4.5	6.2	8.3	2.3	5.0	8.7	7.8
Health	9.8	2.7	6.5	11.8	1.6	6.7	9.2	8.1
Arts	1.2	1.7	1.6	1.2	0.8	1.2	1.1	1.1
Accommodation	4.8	9.8	9.1	5.2	5.6	8.3	5.3	5.0
Services	3.9	3.0	3.8	3.9	2.4	3.9	4.4	4.3
Public Admin.	8.3	8.4	7.9	8.6	3.0	8.4	9.9	11.6
Obs.	9,028,875	14,406	17,034	10,053,946	52,611	33,837	1,697,408	2,116,328

Notes: Each entry shows the percent of workers in each industry in the indicated year and provincial grouping for Stayers, Commuters and Migrants, and 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 or 2012. Commuting and Migrant proportions refer to commuting or migrating to ER provinces.

Table 2: Growth in adjusted log earnings, 2000-2012, Selected Industries

Industry	(1) Including skill matrix	(2) No skill matrix
Oil and gas	0.17	0.24
Mining	0.14	0.23
Utilities	0.1	0.25
Bldg Construction	0.02	0.14
Heavy Construction	0.07	0.2
Textile Products	-0.05	-0.11
Wood Products	-0.07	-0.18
Petroleum Products	0.03	0.02
Chemical Manufacturing	-0.05	-0.09
Computer and Electronics	-0.17	-0.14
Transport Equipment	-0.08	-0.16
Petroleum Product distribution	0.15	0.25
Furniture Stores	-0.1	-0.12
Electronic stores	-0.12	-0.09
Air Transportation	-0.1	-0.18
Pipelines	0.27	0.3
Financial	-0.1	-0.09
Education	-0.06	-0.07
Health	-0.04	-0.05
Recreation	-0.12	-0.13
Accommodation	-0.04	-0.04
Food Retail	-0.03	-0.02
Public admin.	-0.14	0.11

Notes: Changes measured relative to base group: crop production. Column 1 displays industrial premia calculated from log earnings regressions that 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. Column 2 is based on log earnings regressions with full interactions of gender, immigrant status, and a quadratic in age but without the 7x7 matrices.

Table 3: Estimation Results, Canada, Non-Panel Based, Including 7x7 Matrices

	(1) OLS	(2) 2SLS	(3) OLS	(4) 2SLS	(5) 2SLS	(6) OLS
ΔR_{ct}	0.98*	0.86*	1.00*	0.96*	1.12*	
	(0.31)	(0.42)	(0.31)	(0.32)	(0.37)	
ΔX^P_{ct}	-6.84	21.15				
	(7.99)	(29.99)				
ΔX^T_{ct}	2.87*	4.27	2.32	7.78*	9.68*	
	(1.43)	(2.82)	(1.30)	(2.03)	(2.52)	
ΔEmp_{ct}	0.42*	0.90*	0.42*	0.42*		
	(0.08)	(0.32)	(0.08)	(0.08)		
IV^X_{ct}						0.034*
						(0.007)
Industry × Year	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	14,247	14,247	14,247	14,247	14,247	14,247
R^2	0.41		0.41			0.40
Instrument set		All		$IV1, IV2$ $IV3, IV^X$	$IV1, IV2$ IV^X	
First-Stage:						
SW p-value:						
ΔR_{ct}		0.000		0.000	0.000	
ΔX^P_{ct}		0.029		-	-	
ΔX^T_{ct}		0.000		0.000	0.000	
ΔEmp_{ct}		0.001		-	-	
Over-id. p-value		0.244		0.963	0.430	

Notes: 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 regions by 102 industry cells in four year differences. The dependent variable is the change in regression adjusted log earnings, and cells with less than 20 observations are excluded. ΔR_{ct} is average rent in the economic region, ΔX^P_{ct} and ΔX^T_{ct} are the change in the value of the permanent move and commuting option, respectively. Columns 1, 3, and 6 are estimated via Ordinary Least Squares and columns 2, 4, and 5 are estimated via Two Stage Least Squares. All estimates are unweighted. The instrument set in column 2 includes $IV1, IV2, IV3$ and IV^X, in addition to two instruments discussed in footnote 24. The bottom panel of the table contains the Sanderson-Windmeijer (SW) first-stage statistics for the indicated excluded variables in the 2SLS procedure. The last row shows the p-value for the Hansen J overidentification test.

Table 4: Actual and Predicted Wage Changes: 2000 to 2012

	Actual (1)	Predicted (2)	Proportion (3)
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

Notes: This table contains actual and predicted wage changes for the specified location and demographic group for the period from 2000 to 2012. Column 1 shows the proportional changes in the mean wage. Column 2 shows the predicted change based on the commuting coefficient in column 4 of Table 3, combined with the value of ΔX_{ct}^{T} for the specified group for the 2002 to 2012 period. Column 3 shows the proportion of the mean wage change in column 1 that is predicted predicted in column 2.

Table 5: Estimation Results, Alternative Specifications

	Non-Panel, No Skill Matrix			Panel, with Skill Matrix		
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	2SLS	2SLS	OLS	2SLS	2SLS
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_{ct}^T	0.35	3.29*	3.82*	-0.14	3.33*	4.24*
	(0.57)	(1.31)	(1.48)	(0.56)	(1.39)	(1.65)
ΔEmp_{ct}	0.35*	0.33*		0.34*	0.33*	
	(0.07)	(0.07)		(0.09)	(0.09)	
Industry×Year	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	14,247	14,247	14,247	12,003	12,003	12,003
R^2	0.20			0.99		
Instrument set		$IV1, IV2, IV^X$			$IV1, IV2, IV^X$	
First-Stage:						
SW p-values:						
ΔR_{ct}		0.000	0.000		0.000	0.000
ΔX_{ct}^T		0.000	0.000		0.000	0.000
Over-id. p-value		0.674	0.870		0.898	0.727

Notes: 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 regions by 102 industry cells in four year differences. The dependent variable is the change in regression adjusted log earnings, and cells with less than 20 observations are excluded. In columns 1-3, the regression adjustment does not include the 7×7 skill matrix. In columns 4-6, the regression adjustment includes the 7×7 skill matrix and focuses on a panel of individuals who never move or commute from their economic region over the sample period. ΔR_{ct} is average rent in the economic region and ΔX_{ct}^T are the change in the value of the commuting option. Columns 1 and 4 are estimated via Ordinary Least Squares and columns 2, 3, 5, and 6 are estimated via Two Stage Least Squares. All estimates are unweighted. The bottom panel of the table contains the Sanderson-Windmeijer (SW) first-stage statistics for the indicated excluded variables in the 2SLS procedure. The last row shows the p-value for the Hansen J overidentification test.

Table 6: Log Wage and Earnings Levels, Canada and the US

	Hourly Wage		Weekly Wage		Annual Earnings			
	CPS	LFS	Census	Census	Census	Census	Tax Data	Tax Data
Year	US	Canada	US	Canada	US	Canada	US	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

Notes: Each entry shows the value of the average log wage or earnings, calculated using the indicated data set, year and wage measure. 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 individuals between 20-54 years of age, working for wages and salary.

Table 7: Estimation Results, Full Country Sample, Including 7×7 Matrices

	(1) OLS	(2) 2SLS	(3) OLS	(4) 2SLS
ΔR_{ct}	2.80*	3.21*	2.63*	2.03*
	(0.36)	(0.42)	(0.40)	(0.50)
ΔEmp_{ct}	0.41*	0.40*	0.42*	0.44*
	(0.10)	(0.10)	(0.10)	(0.09)
ΔER Share$_{ct}$			0.85	7.62*
			(0.99)	(2.40)
Industry×Year	Yes	Yes	Yes	Yes
Obs.	18,568	18,568	18,568	18,568
R^2	0.50		0.50	
Instrument set		$IV1, IV2$		$IV1, IV2$ IV^X
First-Stage:				
SW p-value:				
ΔR_{ct}		0.00		0.00
ΔER Share$_{ct}$				0.00
Over-id. p-value		0.04		0.29

Notes: 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. The dependent variable is the change in regression adjusted log earnings, and cells with less than 20 observations are excluded. ΔR_{ct} is average rent in the economic region. ΔER Share$_{ct}$ is the fraction of employment in the resource extraction sector in the economic region. Columns 1 and 3 are estimated via Ordinary Least Squares and columns 2 and 4 are estimated via Two Stage Least Squares. All estimates are unweighted. The bottom panel of the table contains the Sanderson-Windmeijer (SW) first-stage statistics for the indicated excluded variables in the 2SLS procedure. The last row shows the p-value for the Hansen J overidentification test.

Table 8: Decomposition of Resource Boom Wage Effects in Canada

	Prov.	Source	Effect	Share	Contribution	Cumulative	Proportion
			(a)	(b)	(a)×(b)		$100 \times \frac{(a)\times(b)}{0.076}$
(1)	ER	Shift-Share	0.016	0.16	0.003	0.003	3.36
(2)		Spillover	0.154		0.025	0.027	35.82
(3)	Non-ER	Commuting	0.009		0.008	0.035	46.21
(4)		Labour Supply	0.002	0.84	0.002	0.037	48.80
(5)		Induced Demand	0.001		0.001	0.038	49.91

Notes: Decomposition results discussed in section 7. The decomposition is broken down between ER and non-ER provinces, and contains the contribution from each indicated source. The impact of each source is given by (a) multiplied by the size of the province (b), to give the contribution towards total wage change. The last column shows the proportion of the total average wage change cumulatively explained by each source.

Table 9: Estimation Results: U.S

	(1) OLS	(2) OLS	(3) 2SLS	(4) 2SLS
ΔR_{ct}	0.81*	0.79*	2.14*	1.80*
	(0.093)	(0.093)	(0.45)	(0.51)
ΔX^T_{ct}		2.44*		5.58*
		(0.87)		(2.17)
ΔEmp_{ct}	0.24*	0.24*	0.081	0.17
	(0.051)	(0.051)	(0.22)	(0.22)
Observations	51809	51809	51809	51809
R^2	0.048	0.049	.	.
Industry× Year	Yes	Yes	Yes	Yes
Instrument set:			IV1-IV3	$IV1 - IV3$ IV^X
First-Stage:				
SW p-value:				
ΔR_{ct}			0.00	0.00
ΔEmp_{ct}			0.00	0.00
ΔX^T_{ct}				0.00
Over-id. p-val			0.80	0.70

Notes: Standard errors, in parentheses, are clustered at the commuting zone level. The asterisk (*) denotes significance at the 5% level. All models are estimated using 584 commuting zones by 134 industry cells, using the 2000 Census and the ACS 2005/06, 2009/10, and 2014/15. The dependent variable is the change in regression adjusted city-industry log wage, and cells with less than 20 observations are excluded. Columns 1-2 are estimated via Least Squares and Columns 3-4 are estimated via Two Stage Least Squares. The bottom panel of the table shows the results of the Sanderson-Windmeijer (SW) first-stage statistics for the excluded variables of the 2SLS procedure for columns 3-4. The last row shows the p-value for the Hansen J overidentification test.