March 2017

COLLABORATION FOR APPLIED RESEARCH IN ECONOMICS

ASYMMETRIC DEMAND SHOCKS AND REGIONAL LABOUR MARKET ADJUSTMENTS IN CANADA

Michael Sullivan

Prepared for the Collaboration for Applied Research in Economics (CARE) initiative,
Department of Economics, Memorial University of Newfoundland

MEMORIAL UNIVERSITY

Abstract

This paper provides a descriptive overview of disparities across Canada's provincial labour markets since the 1970s and estimates a panel vector autoregressive model capturing regional adjustment dynamics. The results of this estimation suggest that region-specific, or asymmetric, labour demand shocks have lasting changes on provinces' relative employment levels. These shocks affect employment mainly through a migration channel and a participation channel: Canadians migrate to provinces with relatively strong labour markets and demand shocks have persistent effects on provinces' labour force participation rates. These channels account for quantitatively similar post-shock employment changes. I also discuss both the robustness of this paper's conclusions to changes in its baseline model's specification and a forecasting application of said model.

Acknowledgments

I thank Collaborative Applied Research in Economics at Memorial University of Newfoundland for its financial support of this project, and Doug May for providing thoughtful feedback on its progress.

1 Introduction

Differences in labour market performance across Canada's regions have long attracted attention from pundits, policymakers, and economists. Much of this attention has focused on interprovincial migration; the settling of the West in the nineteenth century, the Anglophone exodus from Quebec in the later twentieth century, and the decline of Newfoundland's population following the collapse of its cod stocks in the 1980s and 1990s provide historical examples of economic differences driving migration. Interest remains in understanding the economic and migration-related consequences of regional shocks; consider, for instance, the Bank of Canada's April 2016 Monetary Policy Report, which claimed that interprovincial migration in response to the past years' oil price shock was reinforcing Canadian house price divergences (Bank of Canada 2016). Understanding this shock and others likely requires an understanding of past asymmetric shocks.

The policy considerations associated with interprovincial migration are manifold; some of these are explicitly noted by researchers studying labour mobility. Courchene (1970) and Fallahi and Rodriguez (2015) note the role of employment insurance in potentially inhibiting migration flows and hence preventing regional macroeconomic adjustment. Brown and Scott (2012) ask why human capital concentrates in certain areas, a question relevant to policymakers and firms aiming to attract productive workers. And, as Coulombe and Tremblay (2009) remark, the outmigration of skilled workers could lower provinces' returns to educational spending and hence prompt underinvestment in schooling.

Past decades have witnessed the development of a literature exploring regional economies in Canada and intranational labour mobility. Courchene (1970) considers the determination of migration rates between pairs of provinces using aggregated data, motivating his study by noting the scale of regional disparities in Canada and asking: "what are the avenues through which regional adjustment proceeds?" He finds evidence that earnings and unemployment rate differentials between provinces influence migration frequency between them in the expected directions. Additionally, Courchene concludes that federal government transfers to provinces inhibit outmigration; that, presumably as a reflection of urbanisation, provinces with larger agricultural populations experienced greater outmigration; and that younger people were more likely to migrate than older people. More recent work by Coulombe (2006) has added to our understanding of aggregate migration rates. This study, conducted using a pooled time-series cross-sectional approach, suggests structural differences between provinces—say, long-standing differences in labour productivity—drive interprovincial migration to a greater extent than asymmetric provincial business cycles. Dubé and Polèse (2016), somewhat conversely, associate population shrinkage in towns affected by Canada's 2007–2009 recession with these towns' employment and unemployment rate rebounds following the recession, suggesting migration as an adjustment to short-run shocks.

The role of intranational migration in facilitating economic adjustment has received attention outside the Canadian context. This paper follows the broader literature studying this matter and, in particular, the extent to which changes in unemployment, participation rates, and population account for employment shocks. Blanchard and Katz

(1992), notably, examine interstate economic dynamics in the United States with a focus on labour mobility. Their model implies that states' labour demand shocks do not affect wages in the long-rug but that they permanently affect population by prompting migration, which is central to states' macroeconomic adjustment. This study also present estimates of a panel vector autoregressive (VAR) model suggesting that migration provides the primary medium-term adjustment mechanism for regional labour demand shocks. Obstfeld and Peri (1998) conduct similar analyses, expanding the VAR approach of Blanchard and Katz (1992 to Canada and selected European countries. They find that, in the subnational economies of the United Kingdom, Germany, and Italy, labour demand shocks are reflected in participation changes more so than in employment changes relative to Canadian provinces and especially to American states. Relatedly, Obstfeld and Peri (1998) conclude that federal government transfers to regions play a larger role in internal adjustment in Canada and European countries than in the United States, where migration plays an outsized role. A more recent paper by Fahri and Werning (2014) constructs a theoretical model of labour mobility within currency zones that describes when migration aids in regional adjustment to asymmetric demand shortfalls.

Other studies have used longitudinal data at the individual level to understand interprovincial migration. Finnie (2001), for instance, uses Statistics Canada's Longitudinal Administrative Databank (LAD) in studying the relationship between interprovincial migration choices and income; he finds that interprovincial migration was associated with significant increases in individuals' earnings, with young men migrating from relatively poor to relative prosperous provinces benefiting most from migration. Bernard (2011) provides a more recent exploration of LAD data and considers mobility across census agglomerations and census metropolitan areas. Chan and Morissette (2016) use another longitudinal dataset, Statistics Canada's Longitudinal Worker File (LWF), in their attempt to determine the effect of annual wages on interprovincial mobility. Other recent studies, including Sharpe et al. (2007) and Coulombe and Tremblay (2009), have considered the characteristics of interprovincial migrants and the effects of mobility on output and Canada's regional distribution of human capital.

Booms and busts disproportionately affecting some regions may induce potentially persistent changes in unemployment and labour force participation in addition to migration. This paper's interest in asymmetric shocks' effects on provinces' relative participation rates relates to studies of hysteresis in participation. Duval et al. (2010), for instance, consider how recessions of varying severity affected participation rates in a panel of 30 countries, and Clark and Summers (1982) emphasize the persistence of labour supply decisions and relate this to the ability of short-run macroeconomic fluctuations to affect medium-term conditions.

Instead of exploring individuals' interprovincial migration behaviour with microdata, this paper describes trends and fluctuations in provinces' aggregate labour market variables in recent decades and constructs a simple VAR model capturing these dynamics. Stylized facts from the recent history of provinces' labour markets and the results of estimating this model raise questions concerning their microeconomic underpinnings. One of my chief results, for instance, is that provinces' labour markets respond to asymmetric labour demand shocks through both lasting participation rate changes and through temporary changes in net migration rates; these phenomena ac-

count for similar shares of the permanent employment changes following shocks. This raises questions about institutional factors determining the relative roles of migration and participation rate changes in employment shocks, about what makes workers more likely to move or drop out of the labour market, and about the degree and nature of heterogeneity in responses to employment shocks across industries.

The paper's content is divided into two principal sections. Section 2 describes Canada's history of asymmetric economic shocks in the past few decades. Section 3 then presents a panel VAR model of aggregate labour force variables, describes the data and econometric approach used in estimating this model, and interprets its estimated coefficients. Addtionally, section 3 describes (i) robustness checks to and (ii) forecasting applications of its baseline VAR model.

2 Asymmetric shocks: facts and history

The importance of understanding asymmetries in provincial labour market dynamics depends on the degree to which provinces' labour markets mirror or stray from national dynamics. This section begins developing an understanding of shared and of idiosyncratic labour market dynamics by asking how much of provincial employment changes can be explained by national changes. Following Blanchard and Katz (1992), I estimate the coefficients of the equation:

$$\Delta n_{it} = \alpha_i + \beta_i \Delta N_t + \epsilon_{it}. \tag{1}$$

Here, $\Delta N_t$is the first difference of the logarithm of national employment, Δn_{it} is the first difference of the logarithm of employment in province i,[1] ε_{it} is an error term, and α_i and β_i are model parameters estimated for each province; β_i has a natural interpretation as the elasticity of provincial employment with respect to national employment, which I will call sensitivity. Analogously, I estimate equation (1) with industry employment in the place of provincial employment. My estimates of sensitivity, which are reported by table 1, provide rough measures of the historical importance of common trends in employment across provinces and sectors. This exercise's employment data comes from Statistics Canada's Labour Force Survey (see Appendix A for details).

[1] National employment here is employment in Canada minus employment in province i. The approach taken in the industry-level regressions is analagous; instead of subtracting provinces' employment from total Canadian employment, I subtract industries' employment.

Table 1: Sensitivity of provincial and industrial employment to national changes

Province	**Sensitivity**	**Standard error**	R^2
NL	1.11	0.19	0.48
PEI	0.61	0.15	0.32
Nova Scotia	0.90	0.14	0.54
New Brunswick	0.86	0.11	0.61
Quebec	0.94	0.10	0.70
Ontario	1.01	0.12	0.65
Manitoba	0.53	0.07	0.61
Saskatchewan	0.32	0.14	0.13
Alberta	0.82	0.21	0.29
British Columbia	0.80	0.19	0.33

Industry	**Sensitivity**	**Standard error**	R^2
Manufacturing	2.07	0.41	0.41
Education	-0.01	0.20	0.00
Health care	0.18	0.14	0.04
Professional services	1.43	0.315	0.34
Agriculture	-0.24	0.38	0.01
Retail trade	0.63	0.21	0.26
Wholesale trade	0.79	0.54	0.08
Extractive industries	1.30	1.07	0.05
Transportation	1.27	0.23	0.45
Accommodation, food	1.01	0.26	0.29

Note: sensitivity is the ordinary least squares estimate of β_i from equation (1) for the indicated industry or province i. The "Standard error" column presents these estimates heteroskedasticity-consistent standard errors. This table includes estimates only for ten selected industries. Table 3 provides estimates for other industries. Additionally, table 3 provides a correspondence between my industry labels and NAICS codes.

The results reported by table 1 suggest a high level of geographical and industrial heterogeneity in the sensitivity of employment to national changes. The dynamics of a few industries, namely education, health, agriculture, and extractive industries, seem especially independent of national employment changes. These industries' fortunes may highly depend on government spending and international commodity prices, which may in turn be weakly correlated with national macroeconomic performance. Although much of some industries' employment changes can be explained by other industries' changes, there is no industry whose dynamics can be mostly explained by national dynamics,[2] suggesting that industries' fortunes are largely driven by idiosyncratic factors. Industrial heterogeneity in employment dynamics provides one source of the similar heterogeneity in provincial employment dynamics given provinces' different industry mixes. Irrespective of what other forces may explain provinces' divergences from national trends, table 1 justifies the study of these asymmetries by relaying

[2]That is, the R^2 of each industry-level sensitivity regression is below 0.5.

the limits of common trends in explaining employment developments in provinces.

Next, I briefly consider the recent history of labour market trends in Canada's provinces as reflected in aggregate data. Figure A.2 provides us with a few stylized facts about the interaction between labour market variables in Canada's provinces using employment, participation, and unemployment data from Statistics Canada's Labour Force Survey and wage data from its Survey of Employment, Payrolls, and Hours.[3] A plot comparing contemporaneous changes in employment and unemployment, Figure A.2, shows that these variables are inversely correlated. This observation is compatible with an interpretation of employment shocks as labour demand innovations. Indeed, if employment shocks were consequences of supply increases we could expect both employment and unemployment to rise simultaneously, or at least for unemployment to remain roughly constant. Blanchard and Katz (1992) also claim to have consistently found that positive employment shocks increased wages. This relationship is not evident in a first naive view of the Canadian data; as Figure A.2 shows, employment and real hourly wage increases are almost entirely uncorrelated in our sample.

Figure A.2 provides, first, a brief history of real earnings dynamics by province in Canada. Although provincial earnings trends generally have a similar shape as the national trend, these data suggest persistent gaps in wages between provinces, with earnings levels in Maritime provinces and Manitoba remaining well below the Canadian average throughout the entire sample period. Unemployment rates in the provinces, reported by Figure A.2, seem to tend back to national levels following divergences with some notable exceptions, perhaps reflecting the finding of Fallahi and Rodriguez (2015) that regional unemployment rates converge toward the national one, with migration potentially playing a role in this adjustment. Nonetheless, some secular divergences in provincial unemployment rates appear in the data. The unemployment rates in the Atlantic provinces—and especially Newfoundland and Labrador and Prince Edward Island—have trended above national levels throughout our sample, whereas the Prairie provinces have typically boasted lower-than-national unemployment rates. Figure A.2 shows that secular differences are much more pronounced in provincial participation rates. Consistent with the results of unit root tests of my participation rate panel, whose results are presented by table 4, it seems difficult to reject from a consideration of this plot that changes in participation rates dissipate over time at all; this suggests, in turn, that provincial labour markets could adjust to shocks through their denizens' decisions to indefinitely enter or leave their labour forces.

Figure A.2 presents our panel of provincial labour force variables in different manner by shows their cumulative growth in relative to their national levels.[4] Considering this plot, we may first notice that the cumulative employment rate series tend to stay close to zero; in other words, unemployment rates tend to revert to a provincial mean rather than experience secular growth or decline relative to national changes. Conversely, changes in participation rates and especially in employment seem to accumulate over time, showing little regression back to zero. A few other features of these plots are particularly notable. In Alberta, employment grew prodigiously throughout most of the sample. An exception occurred in the 1980s, when the province's oil in-

[3] Appendix A describes these data.

[4] These charts are inspired by those presented by Arpaia et al. (2014).

dustry saw its fortunes momentarily reverse themselves. Although not readily obvious given the Alberta chart's vertical scaling, unemployment more than tripled from 1980 to 1983, and the unemployment rate did not return to its pre-bust levels until the mid-2000s. The participation rate also stopped its ticking upwards, but remained relatively high. Despite this shock, the labour force failed to stop its consistent growth. Alberta's economic turmoil, then, seems to have mostly reflected itself in higher unemployment and slower participation rate growth rather than a migratory adjustment.

Now consider a qualitatively different adverse labour demand shock: that associated with the collapse of Newfoundland and Labrador's cod fishery in the late 1980s and early 1990s. Figure A.2 evinces the drastic decline in the province's relative employment during that period. Although the share of labour force participants who were employed did grow less than the national rate, we did not see the unemployment rate surge above historical levels; it remained largely below the levels of the mid-1980s, in fact. Neither did participation in the labour force decline – conversely, it experienced some of the fastest growth in the country. Instead, outmigration plausibly prevented the province's already underperforming labour markets from spiraling into even direr territory, and Newfoundland and Labrador's population contracted by about 10 per cent from 1991 to 2001. Newfoundland and Labrador's economy experienced a short-lived boom from about 2008 to 2011. Although the unemployment and participation rates had not been responsive to the employment shock during the province's prior bust, they reacted in the boom: the number of unemployed people reached a new minimum in our sample, and the participation rate reached its peak since 1972 in 2011 while the national participation rate began a decline on which it has been ever since.

It may seem surprising that Figure A.2 suggests that the Atlantic provinces experienced participation rate growth far outpacing the national levels whereas Ontario's participation rate growth was particularly slow. Indeed, it is a popular narrative that demographic shifts lowering participation rates are most severe in the Atlantic region. But the high growth of Atlantic region participation rates, instead, mostly reflects that these rates were low at the sample's beginning and trended toward national levels throughout. Figure A.2 shows that this gap-closing process has not even been occurring quickly.

3 A panel vector autoregressive model of provincial labour markets

3.1 Model, data, and estimation

As argued by Sims (1980), vector autoregressions (VARs) provide a flexible framework for capturing macroeconomic variables' interdependencies without relying on unverified exclusion restrictions. This section presents a panel VAR approach for understanding asymmetric shocks in Canadian provinces. This model incorporates five variables that have been manipulated to represent their divergence in each province from their national level. The baseline system we estimate, which is inspired by that presented in Blanchard and Katz (1992) may be schematically written as follows:

$$\begin{bmatrix} \text{average hourly wage (first difference)} \\ \text{employment (first difference)} \\ \text{hours (first difference)} \\ \text{unemployment rate} \\ \text{participation rate (first difference)} \end{bmatrix}.$$

Each of these variables enters the model as the log of its ratio to its national level. In the case of variables that enter the model in first differences of their log relative levels, their transformed values approximate their level percentage changes at the provincial level.[5] I included provincial fixed effects and the first two lags of each variable in estimating each of the model's constituent equations by ordinary least squares (OLS). In computing orthogonalized impulse response functions, I use the LDL decomposition of the OLS covariance matrix as suggested by Hamilton (1994); hence the ordering of variables in the VAR reflect structural assumptions concerning their contemporaneous dynamics. Additionally, I use a nonparametric bootstrapping procedure to estimate the standard errors of the model's orthogonalized impulse response functions.[6]

I use aggregate data from Statistics Canada in estimating the models discussed in this section; these data are described in Appendix A. To summarize, my employment, labour force, unemployment rate, participation rate, and hours series for Canada and the provinces were drawn from the agency's Labour Force Survey. My wage series (which reports average hourly wage for workers paid by the hour including overtime and aggregated over industries) comes from Statistics Canada's Survey of Employment, Payrolls, and Hours.

My specification can be justified by an appeal to stylized facts. In the model, each variable's contemporary value is determined by two lags of each variable and potentially other variables' contemporary values. We must place restrictions on our variables' contemporaneous interdependencies, though, to identify our variables' shocks and trace out their effects. One assumption underlying these restrictions is that wages may not be contemporaneously influenced by changes in employment, the unemployment rate, or the participation rate; I motivate this assumption by appeal to the lack of an empirical correlation between wage changes and employment in my sample.[7] Second, I follow Blanchard and Katz (1992) in interpreting employment innovations as labour demand shocks. Figure A.2, in reporting a negative relationship between provincial shifts in employment and unemployment, suggests the plausibility of this assumption.

[5] Consider employment's representation in the model, for instance. Let us represent provincial employment as e_t and national employment as E_t while defining employment's transformed value in the model as $\delta_t^e = \log(e_t/E_t) - \log(e_{t-1}/E_{t-1})$. When national employment remains roughly constant, we can make the approximation $\delta_t^e \approx \log(e_t/e_{t-1})$. When we represent the percentage change in e_t as $\Delta_t = e_t - e_{t-1}/e_{t-1}$, our approximation becomes $\delta_t^e \approx \log(\Delta_t + 1)$. The first-order Taylor approximation of $\log(\Delta_t + 1)$ is just Δ_t, so we can view δ_t^e as an approximation of Δ_t when national employment remains roughly constant and Δ_t is small.

[6] In this procedure, I simulated samples by drawing from the sample of residuals corresponding to the main model estimation with replacement and then estimated impulse response functions from our simulated data. After estimating the model on 100 simulated samples, I took the empirical standard deviations of the coefficients' empirical distributions as their standard errors.

[7] Figure A.2 plots the correlation between employment changes and wages changes, between which the correlation coefficient is 0.028.

Unlike Blanchard and Katz (1992), I include the relative participation rate in my model in first differences rather than in levels. This reflects the results of panel unit root testing, which suggests that relative participation rates in Canada are nonstationary.[8] The persistence of changes in participation rate changes suggests that economic events affecting this rate—such as business cycles or policy changes—have lasting effects on provinces' labour markets. Numerous studies have considered whether and how business cycles have enduring effects on participation.[9] The VAR modelling approach in this paper adds to this literature by presenting a framework for estimating the extent of participation rate hysteresis and the consequences of cyclical booms and downturns for labour force participation at a subnational level.

3.2 Results

Figure 1 reports the estimated impulse response function coefficients and their confidence intervals with radii of two standard deviations. These estimates imply, first, that employment remains above its pre-shock levels in the medium term following a positive employment shock, and although its trajectory following the shock is somewhat variable, the employment level's entire confidence interval remains decidedly above zero. The impulse response functions for participation rates suggest that much of this increase in employment owes to an increase in participation rates, which also remain above their previous levels following a shock. The effect of an employment innovation, which I interpret as an asymmetric labour demand shock, on the unemployment rate seems to be fleeting, decaying rapidly in the years following the shock. This does not suggest that provincial unemployment rates converge to a national rate, but rather that they tend back towards their provinces' normal rates. Additionally, that the persistent increases in employment precipitated by labour demand shocks tend to be larger than those in participation rates suggests that interprovincial migration is also required to explain the employment effects of such innovations.[10]

These narratives, considered together, imply some basic facts about regional economic adjustment in Canada. First, relative labour demand shocks seem to prompt decreases (or, in the case of adverse shocks, increases) in unemployment that are soon after mitigated by in-migration (or out-migration) and rising (or falling) participation. The point estimates of the impulse response functions plotted in Figure 1 suggest that, two years after the shock, the increase in the participation rate accounts for between a third and a half of the total change in employment. As suggested by Figure 1's wages panel, which shows that the response of relative wages to an employment innovation is highly variable, we cannot confidently conclude how demand shocks affect earnings.

[8] I conducted the panel unit root tests proposed by Levin et al. (2002) and Im et al. (2003) for each variable entering my VAR model. See appendix table 4 for the results of some of these tests.

[9] Two such studies are Duval et al. (2010) and Clark and Summers (1982).

[10] Holding the working-age population constant, an $x\%$ increase in the participation rate corresponds to an $x\%$ increase in the labour force. Given that employment equals $L \times (1 - r^u)$, where L is the labour force and r^u is the unemployment rate, our $x\%$ increase in the participation rate corresponds to an $x\%$ increase in employment when the population and unemployment rate are constant. If employment has increased by more than the participation rate, then either the population or the unemployment rate must also have changed.

These findings suggest a decomposition of employment changes in a province following a labour demand shock. Impulse response coefficients suggest that three years after a positive one per cent employment innovation, the consequent unemployment rate decrease accounts for a 0.17 per cent increase in employment, and the associated participation rate increase accounts for a 0.59 per cent rise in employment. Taken these increases as given, this implies that a population expansion rounds out these other effects by accounting for a 0.71 per cent increase in employment levels. Hence participation rate changes and migration, our assumed cause of population growth, account for the bulk of employment changes following shocks in similar magnitudes, whereas unemployment responses play a smaller role in adjustment.

Although the participation rate enters my primary model in first differences, I find that models in which it enters in its relative-to-national level also suggest that changes in participation account for much of the medium-term employment effects of labour demand shocks. Figure A.2 shows the employment-shock impulse response functions for a model mirroring my preferred specification except that participation enters the model in levels rather than in first differences; this Figure also suggests that participation rates do not return to their pre-innovation levels following labour demand shocks.[11] This finding runs counter to that of Blanchard and Katz (1992) in earlier United States data and somewhat fits with Obstfeld and Peri's (1998) narrative that, compared to the United States, unemployment and participation changes play a larger role in Canada and especially in Europe in adjustment following labour demand shocks.

Additionally, an attempt to understand the dynamics of interprovincial migration through provinces' employment, participation rates, and unemployment rates may strike one as unnecessarily circuitous when provinces' net interprovincial migration rates could be included in the VAR model directly. I estimate such a model and present its impulse response functions in Figure A.2. As this plot shows, explicitly adding migration to the VAR does not significantly alter the responses of the other principal variables included in the model and corroborates the conclusion that migration plays a significant role in provinces' macroeconomic adjustment following shocks. In fact, the impulse response coefficients plotted in Figure A.2 imply that a labour demand shock increasing employment by one per cent raises a province's population by 0.3 per cent by the third year after the shock.

[11] It does, however, show that participation rates exhibit some slow tendency back to their pre-shock levels in this alternative model.

Figure 1: Responses to a labour demand shock

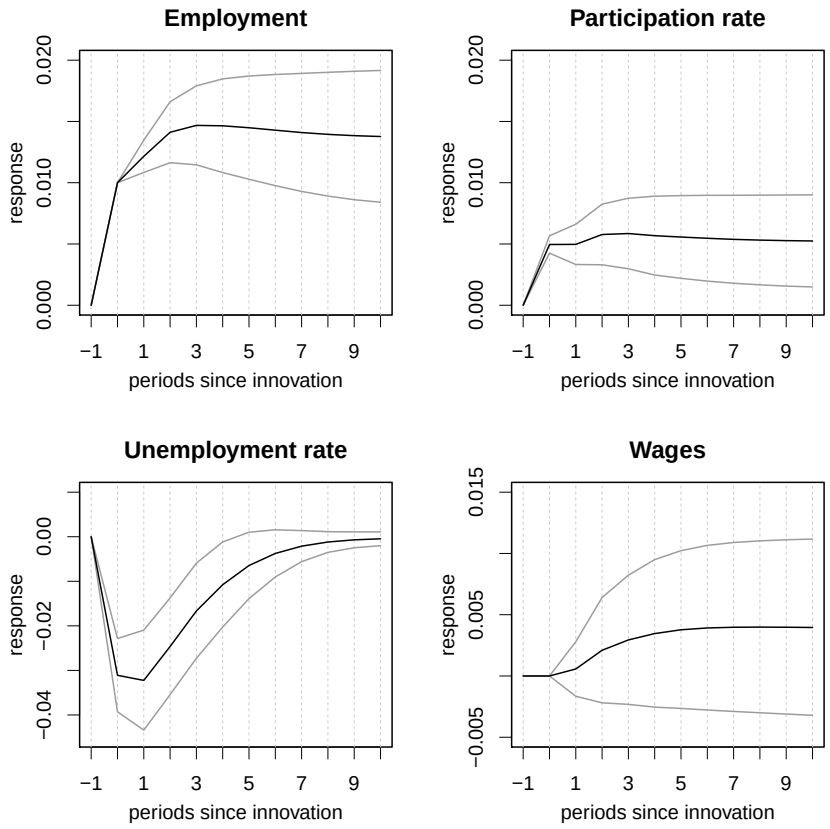

Note: The orthogonalized impulse response coefficients plotted here correspond to a positive employment innovation increasing employment levels by one per cent. Responses are given in percent changes from to the base year (or, one period before the innovation).The light lines above and below the dark line in each panel bound intervals of ± 2 standard deviations of each estimated impulse response function.

3.3 Forecasting and model evaluation

The estimated coefficients of my VAR model imply forecasts for changes in provinces' relative employment, unemployment rates, and participation rates. Table 2 reports these forecasts for the years 2016 and 2017.

As an additional evaluation of our VAR model, we estimate the model on subsets of our data and compute measures of its ability to predict future movements in the variables of interest relative to simple autoregressive (AR) models. Some of these results are provided in appendix table 5. In forecasting employment and participation rates, our preferred model has marginally lower root-mean-square errors (RMSEs) for

two years ahead than those corresponding to forecasts produced using estimates of autoregressive models estimated independently for each variable (the AR forecasts beat the preferred model in predicting unemployment rates). Additionally, models with additional or fewer endogenous variables and different numbers of lags have RMSEs similar to those of the baseline model.[12]

Table 2: Forecasts for relative labour market indicators implied by baseline model

	Employment		**Unemployment rate**		**Participation rate**	
Province	*2016*	*2017*	*2016*	*2017*	*2016*	*2017*
NL	-1.5%	-3.0%	4.8%	6.7%	0.4%	0.4%
PEI	-0.1%	-0.1%	3.8%	2.7%	0.8%	0.8%
Nova Scotia	-0.3%	-0.6%	-0.5%	-0.1%	0.3%	0.6%
New Brunswick	-0.6%	-1.0%	-0.3%	-0.3%	0.4%	0.6%
Quebec	-0.6%	-1.0%	0.9%	2.8%	0.0%	0.1%
Ontario	0.3%	0.6%	-3.8%	-5.5%	-0.1%	-0.1%
Manitoba	0.0%	-0.2%	-0.5%	-2.6%	0.2%	0.2%
Saskatchewan	-0.6%	-1.5%	8.9%	9.9	0.1%	0.0%
Alberta	0.5%	0.8%	6.0%	1.4%	0.1%	-0.5%
British Columbia	0.2%	0.3%	-0.1%	2.5%	-0.3%	-0.4%

Note: the forecasts reported by these table are with respect to the base year of 2015. Percentage changes in indicators' values relative to national levels, not percentage point changes, are reported for the rate variables and for employment. These forecasts were produced using the estimated coefficients of my baseline model.

The results of this forecasting exercise suggest that economies adversely affected by recent collapses in resource prices, especially Newfoundland and Labrador and Saskatchewan, can expect to experience falling employment and increasing unemployement rates in coming years. Conversely, Ontario, whose industry mix is dominated by services-producing, manufacturing, and construction sectors, is forecasted to experience relative employment increases and unemployment decreases, which aligns with our expectations of regional dynamics in the fall-out of an adverse shock to extractive industries located in specific provinces. Some of the forecasts reported by table 2 are more surprising. Take, for instance, that estimates of my baseline model's coefficients suggest relative growth in Newfoundland and Labrador's and Saskatchewan's participation rates and relative decline in those of Ontario and British Columbia. Alberta's employment growth, which stands in contrast to the relative employment decreases of the other provinces with large resource sectors, provides another surprise. This employment increase owes primarily to Alberta's employment growth in 2014 and 2015 and the fact that the estimated VAR coefficients predict, with all else equal, positive contemporary employment changes given positive changes in past years. This illustrates that the infrequency of observations in this model (i.e. each observation considered is a whole year) and its reliance on only a few important economic indicators

[12] Increasing or decreasing the number of lags by one increases RMSE, and we cannot improve our forecasts for each of employment, unemployment rates, and participation rates in an RMSE-sense by switching to one of the other sets of endogenous variables.

could limit the estimates' ability to reasonably predict future outcomes. Interpreters of these forecasts, then, should consider the additional information available to them when predicting future outcomes.

3.4 Robustness

The choices I made in specifying the baseline VAR model discussed above were guided by both empirical and theoretical considerations; this involved considering the sensitivity of my central findings to small changes in the specification. First, I altered the ordering of the endogenous variables in the VAR to determine whether my baseline estimates are sensitive to changes in assumptions about the variables' contemporaneous interactions. I also evaluated my model's robustness to small changes by considering models with more and with fewer endogenous variables than the baseline model (a model explicitly including a migration rate variable, which was discussed above and whose impulse response functions are plotted in Figure A.2, is one example). Estimating these models yielded results corroborating the same qualitative conclusions that I have drawn from the baseline model throughout this section: that employment shocks spur lasting changes in both employment and participation rates; that the greater magnitude of employment shifts imply that interprovincial migration plays a significant role in interregional economic adjustment; and that provinces' relative unemployment rates tend to recover to previous levels fairly quickly after employment shocks.

The ordering of variables in the panel VAR model described throughout this section is, to some extent, arbitrary. In ensuring that my primary conclusions do not depend, I estimated the model and computed its impulse response coefficients for an employment shock under alternative orderings. Figure A.2 presents such impulse response coefficients for a model with an alternative ordering; by comparison with Figure 1, it is clear that my substantive conclusions still hold under this significant reordering.[13] The sole major change evident in this plot, which describes an estimated model in which employment is permitted to contemporaneously affect wages, is that the estimated effect of the employment shock on wages is larger than in my baseline specification.

4 Conclusion

The behaviour of Canada's provincial and industrial economies cannot easily be predicted by national data alone. The idiosyncratic dynamics of its regional economies naturally raises questions about how provinces adjust to positive and adverse demand shocks that are not felt by other provinces. This paper specified a panel VAR model describing the interdependencies of key labour market indicators across provinces, concluding that these shocks have persistent effects on provinces' populations and labour force participation rates as workers leave or join their labour markets in response to changed aggregate conditions. Quantitatively, the number of workers leaving for other provinces and other dropping out of the provincial labour force are similar; shocks also have less profound effects on provincial unemployment rates in the short- and midterm.

[13] See Figure A.2's notes for a description of this alternative ordering.

References

Arpaia, Alfonso, Aron Kiss, Balazs Palvolgyi, and Alessandro Turrini. 2014. Labour mobility and labour market adjustment in the EU. European Economy. European Commission *Economic Papers* 539.

Bank of Canada. 2016. *Monetary Policy Report* (April). Accessed at <http://www.bankofcanada.ca/wp-content/uploads/2016/04/mpr-2016-04-13.pdf>

Bernard, André. 2011. Regional economic shocks and migration. Statistics Canada Perspectives on Labour and Income Catalogue no. 75-001-X.

Blanchard, Olivier J. and Lawrence F. Katz. 1992. Regional evoluations. *Brookings Papers on Economic Activity* 1: 1–75.

Chan, Ping Ching Winnie and Rene Morisette. 2016. The impact of annual wages on interprovincial mobility, interprovincial employment, and job vacancies. Statistics Canada Analytical Studies Branch Research Paper Series 11F0019M No. 376.

Clark, Kim and Lawrence Summers. 1982. Labour force participation: timing and persistence. *Review of Economic Studies* 49(5): 825–844.

Coulombe, Serge. 2006. Internal migration, asymmetric shocks, and interprovincial economic adjustments in Canada. *International Regional Science Review* 29(2): 199223.

Coulombe, Serge and Jean-François Tremblay. 2009. Migration and skills disparities across the Canadian provincies. *Regional Studies* 43(1): 5–18.

Courchene, Thomas J. 1970. Interprovincial migration and economic adjustment. *The Canadian Journal of Economics* 3(4): 550–576.

Dubé, Jean and Mario Polèse. 2016. Resilience revisited: assessing the impact of the 2007-09 recession on 83 Canadian regions with accompanying thoughts on an elusive concepts. *Regional Studies* 50(4): 615–628.

Duval, Romain, Mehmet Eris, and Davide Furceri. 2010. Labour force participation hysteresis in industrial countries: evidence and causes. OECD Working Paper.

Fahri, Emmanuel and Iván Werning. Labor Mobility within Currency Unions. Working paper. Accessed on March 17, 2017 at <http://scholar.harvard.edu/files/farhi/files/mobility.pdf>.

Fallahi, Firouz and Gabriel Rodriguez. 2015. Structural breaks and labour market disparities in the Canadian provinces. *Journal of Economic Studies* 42(2): 322–342.

Finnie, Ross. 2001. The effects of inter-provincial mobility on individuals' earnings: panel model estimates for Canada. Statistics Canada Analytical Studies Branch Research Paper Series 11F0019MIE No. 163.

Hamilton, James Douglas. 1994. *Time Series Analysis*. Princeton, New Jersey: Princeton University Press.

Im, Kyung So, M. Hashem Pesaran, and Yongcheol Shin. 2003. Testing for unit roots in heterogenous panels. *Journal of Econometrics* 115(1): 53–74.

Levin, Andrew, Chien-Fu Lin, and Chia-Shang Chu. 2002. Unit root tests in panel data: asymptotic and finite-sample properties. *Journal of Econometrics* 108(1): 1–24.

Obstfeld, Maurice and Giovanni Peri. 1998. Regional nonadjustment and fiscal policy. *Economic Policy* 13(26): 205–259.

Sharpe, Andrew, Jean-Francois Arsenault, and Daniel Ershov. 2007. The impact of interprovincial migration on aggregate output and labour productivity in Canada, 1987–2006. Centre for the Study of Living Standards Research Report 2007–02.

Sims, Christopher. 1980. Macroeconomics and reality. *Econometrica* 48(1): 1–48.

Appendices

A Data

Sections 2 and 3, and the tables and figures they reference, use aggregate data from Statistics Canada's CANSIM databank. My employment, labour force, unemployment rate, and participation rate series for Canada and the provinces were drawn from annual Labour Force Survey estimates as reported in CANSIM table 282-0008. These data were available for 1976–2015, and their values are for those of fifteen or more years of age. My hours of work series was drawn from annual Labour Force Survey estimates for of fifteen or more years of age, too, but from CANSIM table 282-0016. This series extends from 1976 to 2015 as well.

My wage series are from annual Survey of Employment, Payrolls, and Hours as reported in CANSIM tables 281-0008 and 281-0030. Specifically, I used a series reporting nominal average hourly wage for workers paid by the hour including overtime and aggregated over industries, excepting unclassified establishments. The series I constructed using these data extends from 1983 to 2015.

I used interprovincial migration data and population estimates in constructing my net migration rate series. My migration data come from CANSIM table 051-0017; the annual values in my series are the sums of quarterly values. My Canadian and provincial population data come from Statistics Canada's population estimates for July 1 of each year as reported in CANSIM table 0051-0001. My net migration rate series extends from 1976 to 2015.

B Plots

Figure A.2: Co-movements of employment and other labour force variables

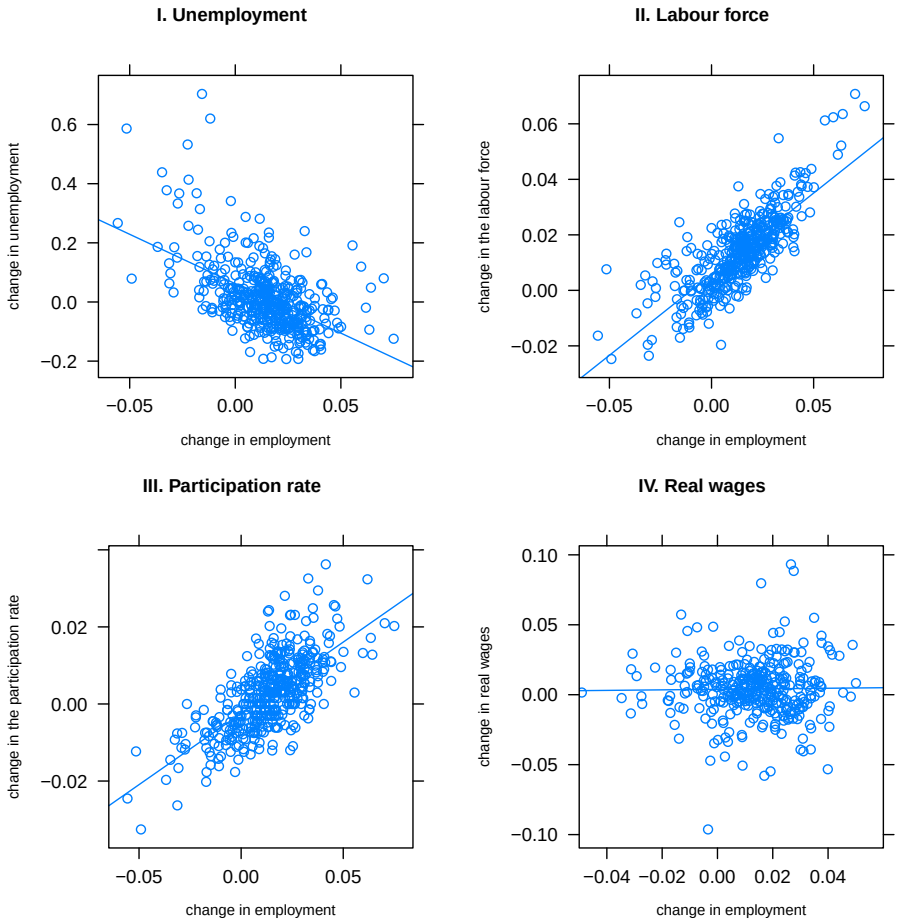

Figure A.2: Real wages across Canadian provinces

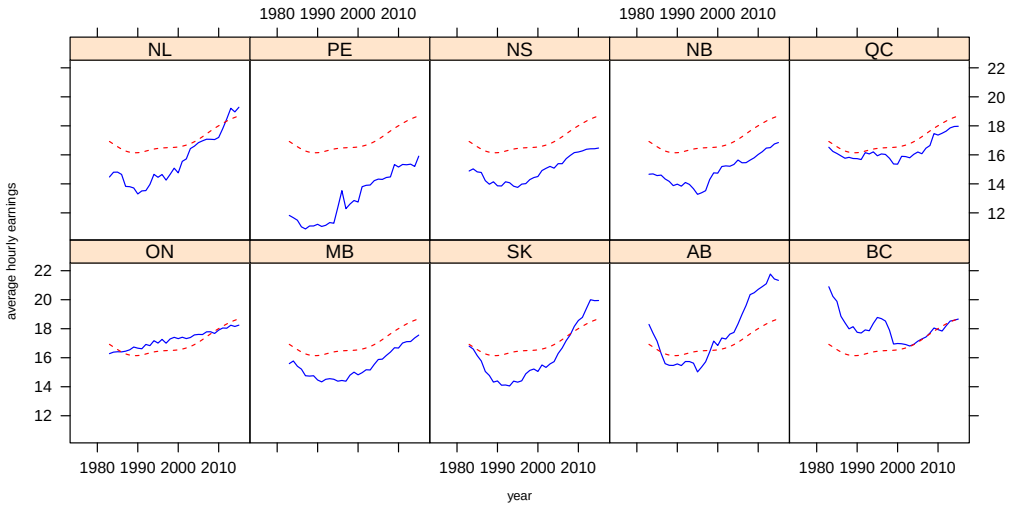

Note: the dashed curve replicated in each plot gives the national hourly wage. Average hourly earnings here are for workers paid by the hour and are expressed in 2002 dollars.

Figure A.2: Unemployment rates across Canadian provinces

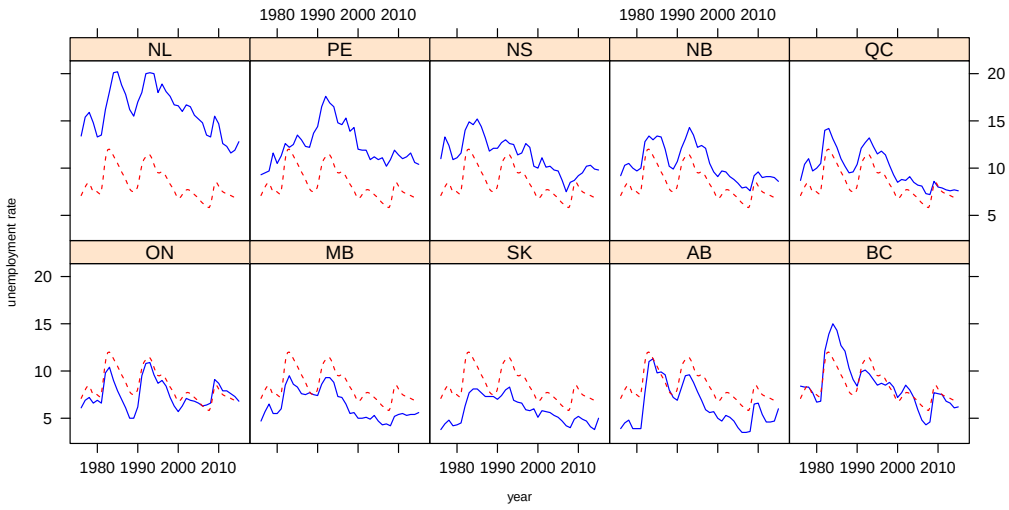

Note: the dashed red curve replicated in each plot gives the national unemployment rate.

Figure A.2: Participation rates across Canadian provinces

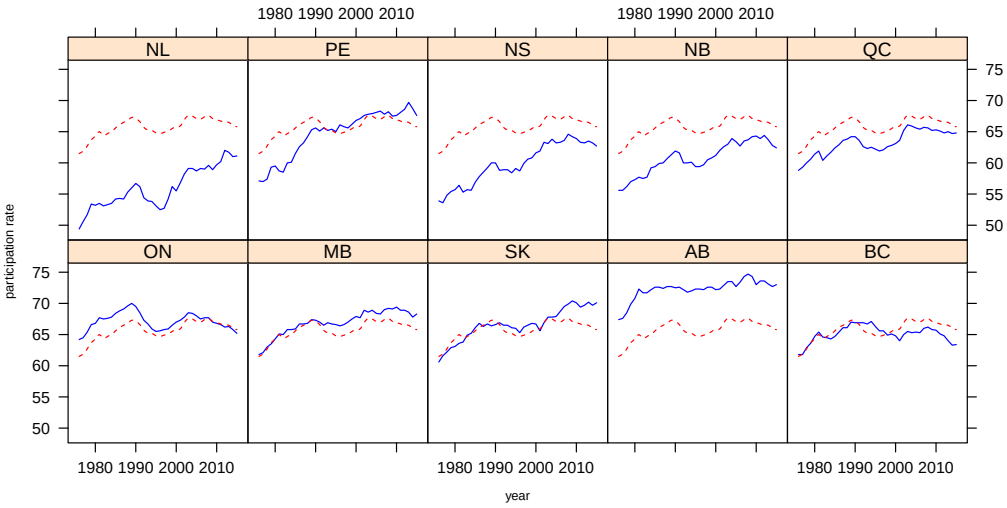

Note: the dashed curve replicated in each plot gives the national participation rate.

Figure A.2: Accumulated growth rates of labour market indicators

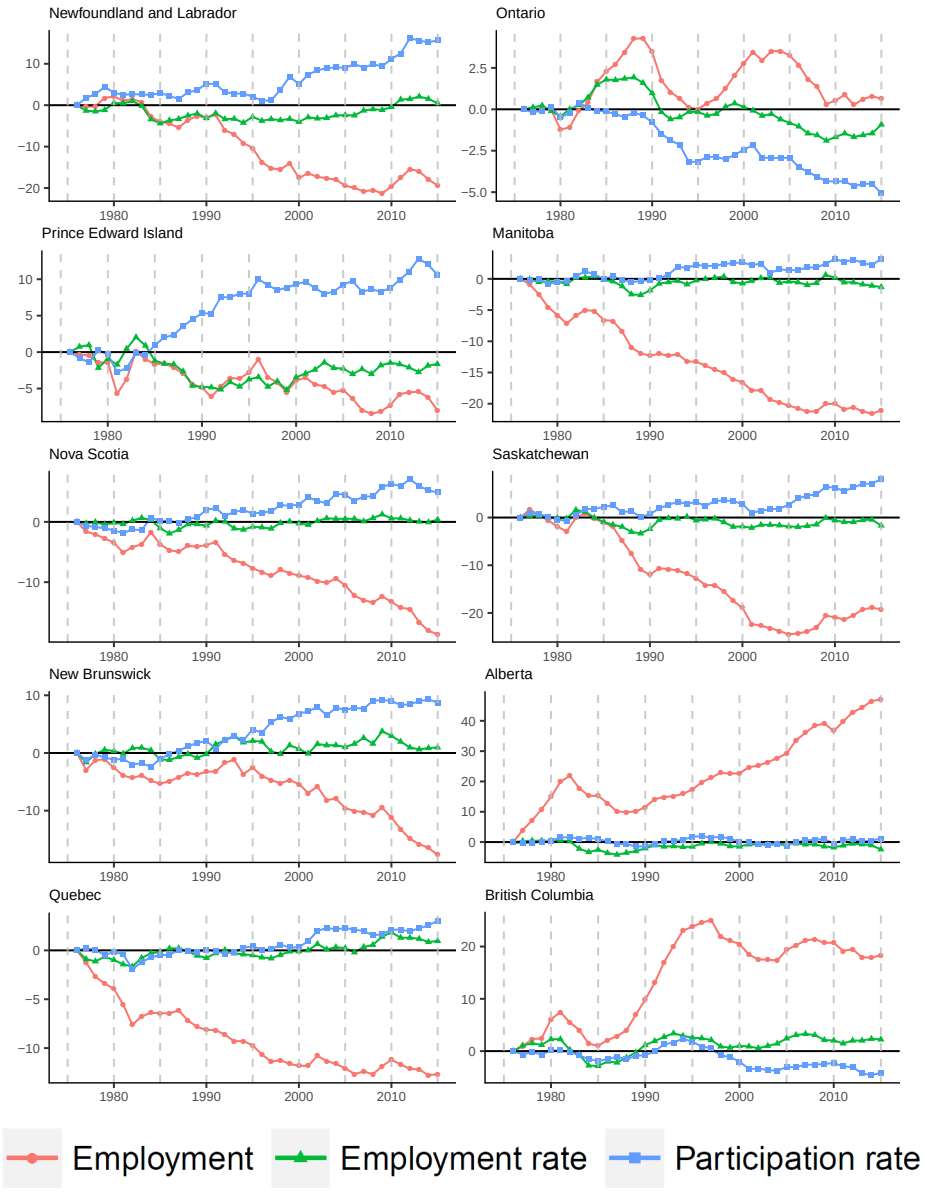

Notes: the employment rate here is defined as one minus the unemployment rate.

Figure A.2: Impulse response functions (participation rate enters model in levels)

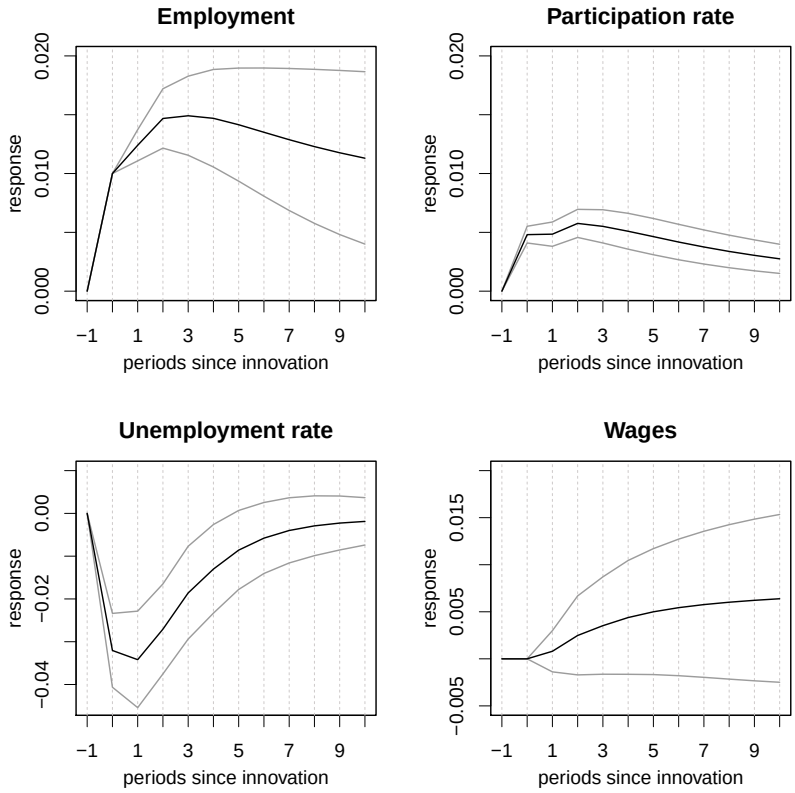

Note: see the notes of Figure 1.

Figure A.2: Impulse response functions (including migration rates)

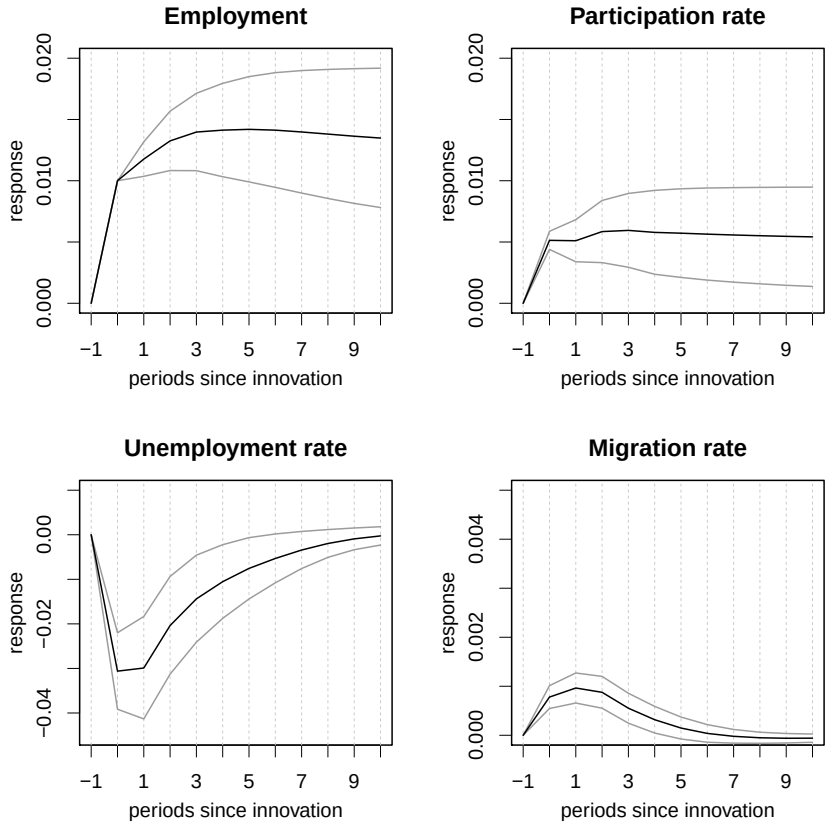

Note: see the notes of Figure 1. The migration rate panel here gives the change of the province's migration rate in levels from a base of zero.

Figure A.2: Impulse response functions (alternative ordering)

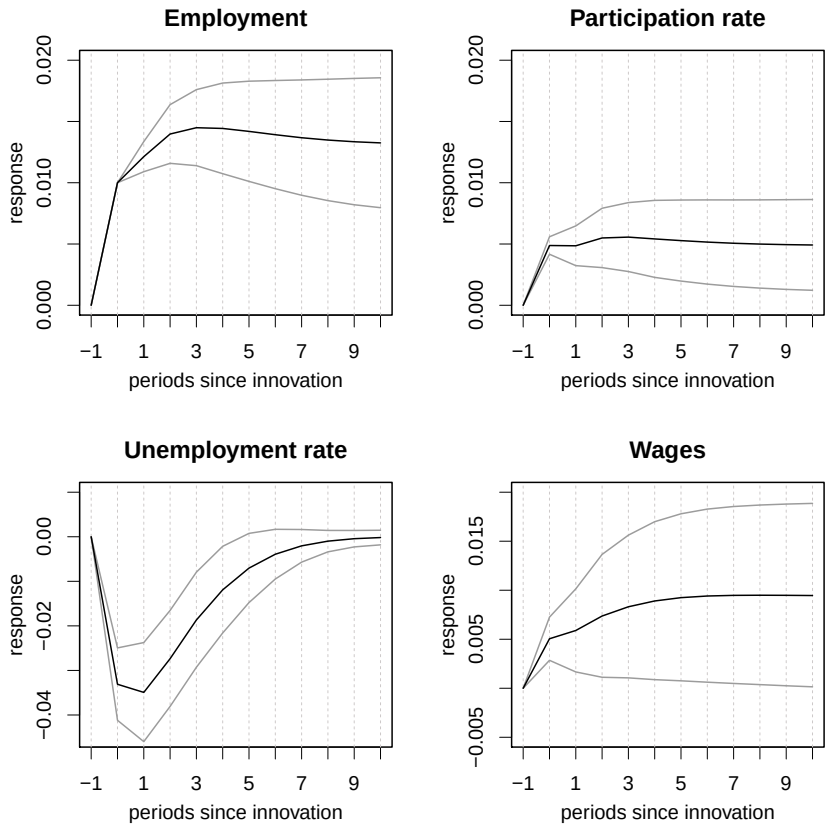

Note: see the notes of Figure 1. The ordering of the endogenous variables in the VAR from whose estimation results I constructed these impulse response functioins is: employment, wages, the participation rate, hours, and the unemployment rate.

C Tables

Table 3: Sensitivity of industrial employment to national employment changes

Industry	**Sensitivity**	**Standard errors**	R^2
Forestry	1.83	1.14	0.09
Fishery	0.44	1.11	0.01
Utilities	0.69	0.53	0.04
Construction	1.68	0.43	0.30
Financial services	0.41	0.38	0.04
Real estate	−0.16	0.64	0.00
Business services	1.09	0.32	0.24
Information, culture, and recreation	0.60	0.33	0.08
Other services	0.46	0.32	0.05
Public administration	−0.02	0.22	0.00

Note: See notes for Figure 1."Forestry" refers to NAICS industries 113 and 1153 (forestry and logging and support activities for forestry). "Fishery" refers to NAICS industry 114 (fishing, hunting, and trapping). "Utilities" refers to NAICS industry 22. "Construction" refers to NAICS industry 23. "Financial services" refers to NAICS industry 52 (finance and insurance). "Real estate" refers to NAICS industry 53 (real estate and renting and leasing). "Business services" refers to NAICS industries 55 and 56 (business, building, and other support services). "Information, culture, and recreation" refers to NAICS industries 51 and 71. "Other services" refers to NAICS industry 81 ("other services, except public administration). "Public administration" refers to NAICS industry 91. From table 1, "Manufacturing" refers to NAICS industries 31, 32, and 33; "Education" refers to NAICS industry 61 (education services); "Health care" refers to NAICS industry 62 ("health care and social assistance"); "Professional services" refers to NAICS industry 54 (professional, scientific, and technical services); "Agriculture" refers to NAICS industries 111, 112, 1100, 1151, and 1152; "Retail trade" refers to NAICS industries 44 and 45; "Wholesale trade" refers to NAICS industry 41; "Extractive industries" refers to NAICS industry 21 (mining, quarrying, and oil and gas extraction); "Transportation" refers to NAICS industries 48 and 49 (transportation and warehousing); and "Accommodation, food" refers to NAICS industry 72 (accommodation and food services).

Table 4: Selected panel unit root test results

Variable	**Test statistics**	
	Levin-Lin-Chu	*Im-Pesaran-Shin*
Employment	-11.67	-11.90
(first difference)	(0.00)	(0.00)
Unemployment rate	-4.41	-4.60
	(0.00)	(0.00)
Participation rate	0.10	1.74
	(0.92)	(0.08)
Participation rate	-16.90	-16.77
(first difference)	(0.00)	(0.00)
Hourly earnings	-11.94	-12.26
(first difference)	(0.00)	(0.00)
Net interprovincial	-6.68	-6.73
migration rate	(0.00)	(0.00)

Note: the statistics' p-values are provided in parentheses. The tests of both Levin et al. (2002) and Im et al. (2003) tests are against a null hypothesis that the panel in question in nonstationary. Individual intercepts were included as the sole exogenous terms in these tests' regression equations. The series are the indicated variables' relative levels or, where identified, their first differences. I transformed the panel of provinces' net interprovincial migration rates as $f(m_t) = \log(1 + m_t)$ *before conducting tests.*

Table 5: VAR model forecast evaluation

Employment

Model	**RMSE**	
	VAR	*AR*
Preferred (one lag)	0.0115	0.0113
Preferred (two lags)	0.0111	0.0114
Preferred (three lags)	0.0116	0.0117
Migration rate included	0.0108	0.0114
Only three variables	0.0116	0.0117

Unemployment rate

Model	**RMSE**	
	VAR	*AR*
Preferred (one lag)	0.107	0.202
Preferred (two lags)	0.110	0.106
Preferred (three lags)	0.117	0.607
Migration rate included	0.111	0.106
Only three variables	0.104	0.103

Participation rate

Model	**RMSE**	
	VAR	*AR*
Preferred (one lag)	0.0073	0.0072
Preferred (two lags)	0.0072	0.0073
Preferred (three lags)	0.0073	0.0071
Migration rate included	0.0073	0.0073
Only three variables	0.0076	0.0075

Note: the "VAR" columns report the root-mean-square errors (RMSEs) produced by estimating VARs on subsets of the sample data and using them to forecast their endogenous variables two years after the sample. The "AR" columns report RMSEs produced using the same method but instead using the results of estimating autoregressive models on subsets of data. These autoregressions featured the number of lags of the dependent variable that were included in the corresponding VAR in addition to provincial fixed effects as explanatory variables. The "Preferred" model is the model proposed by section 3. The "Migration rate included" model is that whose estimated impulse responses functions are plotted in Figure A.2 (this model includes two lags). The "Only three variables" model includes only transformed employment, unemployment rates, and participation rates as explanatory variables (this model includes two lags). This model was estimated on a larger sample than the others as data for its variables were available earlier than for wages, which were included in the other models. The RMSEs reported in this table correspond to the transformed versions of the variables listed. I transform employment and participation rates by taking the first differences of the log ratios of their provincial and national levels. When the national levels are constant, the transformed values are the first-order Taylor approximations of the level variables' percentage changes at the provincial level. I transform unemployment rates by taking the log ratio of their provincial and national levels.