The Impact of Labour Force Participation on Unemployment Rate Fluctuations in Canada

by Graeme E. Byron

(Student #300070685)

Major Paper presented to the Department of Economics of the University of Ottawa in partial fulfillment of the requirements of the M.A. Degree
Supervisor: Professor Jonathan Créchet

ECO 6999

Ottawa, Ontario

August 2024

Abstract

This paper studies cyclical fluctuations in the Canadian unemployment rate using a three-state framework of employment, unemployment and nonparticipation. First, by using a flow-based variance decomposition, I find that transitions between nonparticipation and unemployment can explain around 35% of the change in the Canadian unemployment rate between 1976 and 2019. For comparison, these transitions account for 33% of the change in the American unemployment rate between 1978 and 2012. However, in Canada, nonparticipation to unemployment transitions play a much larger role (21% vs 13%). Additionally, my findings support the conclusions of past papers that a stock-based decomposition of the unemployment rate is subject to the stock flow fallacy. Since over 90 percent of transitions from one labour market state to another gets offset by simultaneous transitions in the opposite direction, a decomposition cannot distinguish between the different patterns of worker flows that result drive cycles in the participation rate. Finally, by comparing the Canadian and American labour markets, I show that during recessions, the percentage of the unemployment population in Canada that belong to each demographic, including categories of age, sex, education level and an individual's reason for unemployment, does not change as substantially as in the United States. This is due to the Canadian labour market being less dynamic.

1. Introduction

For decades, understanding the causes of unemployment rate fluctuations has been a goal of many economists. Current empirical analyses suggest that the participation margin plays an insignificant role in explaining changes in unemployment (Rogerson and Shimer (2011)). However, Elsby, Hobijn and Şahin (2015) challenge this finding with their flow-based decomposition of the unemployment rate using gross flow data in a three-state framework of employment, unemployment and nonparticipation. They find that transitions between unemployment and nonparticipation can account for 33% of the change in the American unemployment rate between 1967 and 2012. In this paper, I use the same decomposition with Canadian gross flow data to determine if the participation margin also contributes significantly to the Canadian labour market. My findings show that the cyclical nature of transitions between unemployment and nonparticipation do have a significant effect on the unemployment rate, accounting for 35% of changes in the unemployment rate between 1976 and 2019. From this result, I conclude that the participation margin should be accounted for in future studies.

Additionally, this paper contributes to the existing literature by providing the first unemployment rate variance decomposition of the Canadian labour market that uses gross flow data as opposed to unemployment duration data.[1] By comparing my results with past papers that have use unemployment duration data (such as Campolieti (2011), Elsby, Hobijn and Şahin (2013)), I show that gross flow data suggest the combined inflow to unemployment from nonparticipation and employment plays an increased role in determining variation of the unemployment rate. Likewise, this phenomenon is observed in studies of the United States labour market (see Shimer (2012), Elsby et al. (2015)).

I further analyze flows in and out of the labour market by using a stock-based decomposition. My findings support the conclusion of Elsby et al. (2015), that this decomposition is subject to a stock flow fallacy. Since over 90 percent of transitions from one labour market state to another gets offset by simultaneous transitions in the opposite direction, a stock-based decomposition cannot distinguish between the different patterns of worker flows that drive cycles in the participation rate.

Finally, I analyze how labour force attachment affects the unemployment rate when economic conditions change. I find that relative to the United States, individuals in the less dynamic Canadian labour market are less likely to enter or leave the labour market throughout recessions. This is shown both in terms of changes in labour market entry and exit rates, and in terms of changes in the demographic composition of the unemployment pool during recessions.

The structure of the paper is as follows. Section two contains a non-exhaustive literature review of previous studies of the unemployment rate. In section three, I describe the data I will use, and potential sources of bias. In section four, I provide my methodology to correct for those biases. In section five and six, I present and analyze the results of my variance decomposition. In section seven I present and analyze the results of my stock-based decomposition. Finally, in section eight, I analyze changes in labour force exit and entry rates during recessions to determine the role labour force attachment plays in the Canadian labour market.

2. Literature review

Economists study the causes of fluctuations in the unemployment rate to better understand the business cycle. Changes in the unemployment rate accounts for a larger proportion of variation in total hours worked, which in turn contributes to the peaks and troughs of the business cycle. Under certain circumstances, it can be in a government's best interest to introduce policies that incentivise an increase or decrease in the unemployment rate in order to reduce recessionary periods or increase expansionary periods

[1] Using gross flow data, or unemployment duration data in an unemployment variance decomposition tend to provide significantly different results. In the US, gross flow data tends to result in a roughly two thirds outflows, and one thirds inflows decomposition of the unemployment rate (see Elsby et al. (2015)), while unemployment duration tends to result in an 85% outflows, 15% inflows decomposition (see Elsby et al. (2013), Shimer (2012), etc…).

of the business cycle. If economist can better understand the causes of fluctuations in the unemployment rate, government can more effectively design their policies to achieve their desired result.

Economist have used theoretical models to analyse the causes of unemployment fluctuations. The Diamond, Mortensen and Pissarides model, developed over a series of papers (Diamond (1982), Pissarides (1985), and Mortensen and Pissarides (1994)), models an economy of workers and firms that are matched to produce an output in the presence of search-and-matching frictions. This model also shows how the unemployment rate can fluctuate based on changes in the job finding and job destruction rate. One shortcoming of this model is that is does not account for labour force non-participation. All workers are either employed or seeking employment.

Krussell, Mukoyama, Rodgerson and Şahin built off the DMP model to develop a three-state general equilibrium model that includes labour force participation (Krussell et al. (2011, 2017, 2020)). One thing they find is that intertemporal substitution has a large role in determining fluctuation in the labour force participation rate. In 'good times' the unemployment pool is increasingly populated by individuals who transition from N to U because it is easier for employees to work their way up the job ladder faster, resulting in higher wage growth. For this reason, it is important to account for transitions between nonparticipation and unemployment when analysing unemployment rate fluctuations.

Many economists have also performed empirical analyses of labour market dynamics using labour force data. For decades, governments have published estimates of measures of gross flows between labour market status/states. In Canada, Statistics Canada began this process in 1983 using the Canadian Labour Force Survey (Fienberg and Stasny (1983)). Although the LFS is not a panel dataset, selected households are surveyed for six consecutive months which allows researchers to observe respondents' transitions between the labour force status categories. Furthermore, they can combine that data with survey weights to estimate the gross flows for the entire economy.

Using gross flows, economist have been able to analyse labour market dynamics. In Canada, Jones (1993) analyses the cyclical components of gross flows, including non-participation, and finds that a strong majority of gross flows can be accounted for by seasonal factors such as the beginning or end of the school year. Furthermore, Jones and Riddell wrote a series of empirical papers (Jones and Riddell (1998, 1999, 2006, 2019)) on how to properly categorize non employed individuals based on behavioural patterns. They argue that in addition to unemployment and nonparticipation, a third category of marginally attached nonemployed individuals exist. An example of marginally attached individuals are discouraged workers.[2] They are not actively seeking employment, and thus act distinctly different than unemployed individuals, but they still desire employment, and thus are more likely to join the labour force in the future, relative to

[2] Discouraged workers are nonemployed working aged individuals who desired employment, but don't try to find a job because they believe there isn't one available for them.

other nonparticipants. Jones and Riddell find evidence of marginally attached individuals in both Canada and the United States whose transition rates into employment are predictably different than either unemployed or nonparticipating individuals.

Campolieti (2011) implements a method developed by Shimer (2012) and Elsby, Michaels and Solon (2009) which generates unemployment inflow and outflow hazard rates measuring the probability respectively of an employed individual becoming unemployed, and an unemployed individual gaining employment. He finds that log changes in the outflow hazard rate accounts for roughly 82% of the change in the log unemployment rates in Canada. One issue with this method is that they assume no workers enter or leave the labour force.

As pointed out in Abowd and Zeller (1985), gross flow data derived from survey data is subject to three main problems, including classification error. They find that in CPS survey data, 12% of individuals who should be classified as unemployed, are wrongly classified as either employed or nonparticipation. This causes overestimates of worker flows as it generates spurious labour market transitions.

In 2015 Elsby, Hobijn and Şahin introduced a new three-state variance decomposition of the unemployment rate that can be used with gross flows. Additionally, they also developed a new technique to test if their findings are robust to classification error corrections which involves examining the effect of recoding gross flows to eliminate sequences of transitions from nonparticipation to unemployment and back to nonparticipation (N-to-U-to-N), as well as U-to-N-to-U transitions. The resulting data produced from this recoding process is referred to as "deNUNified" flows. In their decomposition, they find that transitions between N and U account for roughly 30% the of variance in the US unemployment rate.

As this decomposition is relatively new, it has not yet been widely used to analyse labour markets outside of the United States. This paper is the first to utilize a three-state variance decomposition of the unemployment rate that uses gross flow data to analyse the Canadian labour market. In my paper, I aim to answer two questions. First, what role does the participation margin have in driving fluctuations in the Canadian unemployment rate? Second, how that relationship change during recessions?

3. Data

I utilize individual level confidential-use Labour Force Survey (LFS) data to estimate the gross flows of workers in Canada. I have chosen to use LFS samples from the earliest available month (Jan 1976) until December 2019. During the COVID-19 pandemic, beginning in 2020, the LFS had a sharp increase in their rates of non-response resulting from the suspension of in-person interviews during the COVID-19 pandemic (Brochu and Créchet (2022)). Prior to the COVID-19 pandemic, the LFS non-response rate had not surpassed 14 percent in its 44-year history. Beginning in March 2020, the non-response rate skyrocketed to 22.1 percent and eventually reached 30 percent in September 2020. Although the LFS uses data imputation to replace missing observations, this method becomes more susceptible to bias as the percentage of non-responses increase.

The LFS is fundamentally a series of monthly cross-sectional datasets for a sample of the Canadian population. However, it does utilise a six-month rotating panel design. In any given month there are six rotation groups, and each group stays in the sample for six months before being replaced. The start and end months for each group is staggered such that one of six groups are replaced each month. Trough the use of a unique person identifier variable, researchers can link individuals across time and keep track of changes in labour market status.[3] In any given month, a theoretical maximum of 83.33% of the sample will also be observed in the next month. However, due to factors such as non-responses or individuals moving to another dwelling, I tend to only observe 79.4% of the sample in the following month.

The LFS does not use a purely random sample. Instead, they intentionally over-sample certain sub-groups to ensure in order to have a sufficient sample size to reliably calculate various labour market indicators for less populated regions and demographics that represent a small percentage of the population (Brochu (2021)). Within each desired sub-group, representative households are selected randomly from a cluster of households matching the sub-group criteria, and each member of the household is assigned their own weight. Using respondent's self reported labour force status, and adjusting based on survey weight, I can estimate the stock of employed, unemployed and labour force nonparticipating workers in each month. Similarly, I can exploit the LFS panel design to estimate the gross flows of workers between the three categories in each month.

As pointed out by Abowd and Zellner (1985), survey data is subject to a few drawbacks. The three main sources of error are classification error, time aggregation bias and margin error. I will explain these three issues, and my techniques for addressing them in the methodology section. For now, I can analyze the basic cyclical patterns of worker flows by graphing changes in the monthly transition rates between the three categories over time.[4]

In figure 1, I present the monthly unemployment rate and labour force participation rate throughout my sample. Both graphs contain three lines representing different datasets.[5] All three datasets include the same sample, (noninstitutionalized, non-military Canadian residents 15 years of age and older), however the AZ and DeNUNified lines are adjusted for classification error using different methods that I will outline in section 4. A visual examination of figure 1 provides information on the cyclical nature of each rate in

[3] The identifier variable (HHLDID) is formulated from a series of 10 variables to follow individuals over time.
[4] All monthly transition rates have been adjusted for margin error and are seasonally adjusted using a 12-month weighted moving average. My seasonal adjustments take the average the past six month, the current month and the future six months but the first and last months weighted half as much as the others.
[5] All three datasets include the same sample, (noninstitutionalized, non-military Canadian residents 15 years of age and older), however the AZ and DeNUNified lines have both been adjusted for classification error using different methods that I will discuss in section 4.

Canada. One will notice that that countercyclical unemployment rate rises sharply during recessions, and to a lesser extent, the procyclical labour force participation rate falls during recessions. Additionally, by comparing the unadjusted data, to the AZ and DeNUNified data, one will notice that this pattern is not eliminated or significantly reduced when accounting for classification error.

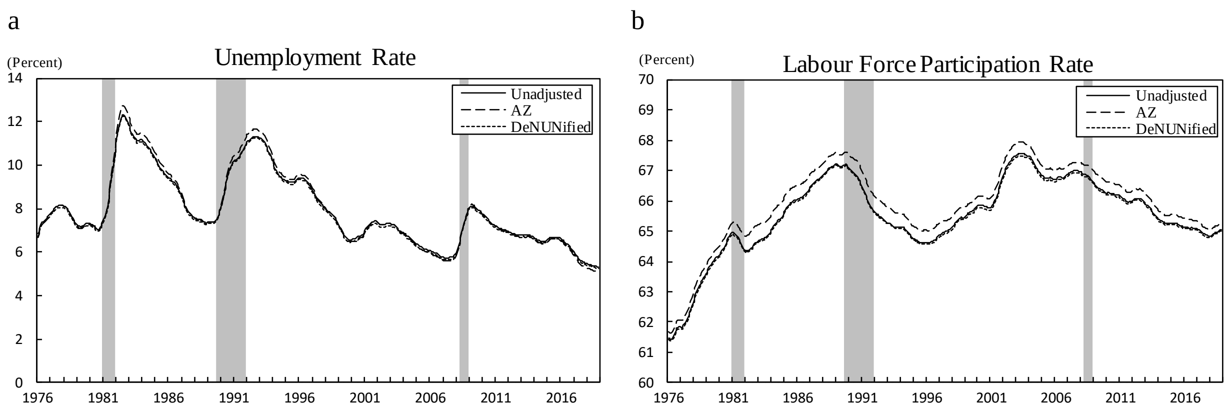

Figure 1: Monthly unemployment and labour force participation rates. These rates are calculated using Canadian LFS data. All stock variables have been seasonally adjusted. The grey bars represent recessions in this sample.

Figure 2 graphs the monthly discrete time transition rates.[6] From a visual examination, one can gain insight about their cyclicality over the business cycle, and how they may impact the unemployment rate. Employment to unemployment transitions is countercyclical and very volatile. Unemployment to employment transitions is similarly volatile but procyclical. Both rates change significantly during recessions contributing to high unemployment rate. Likewise, nonparticipation to unemployment transitions are countercyclical and rise sharply during recessions, however they appear to be slightly less volatile than the two rate previously mentioned. The unemployment to nonparticipation transition rate appears weakly procyclical. Unusually, during the 1990-1992 recession, you do not see a consistent decrease as is seen during the other two recessions. The nonparticipation to employment transition rate is procyclical and consistently decreases during recessions. Finally, employment to nonparticipation is less volatile and should have a limited affect on the volatility of the unemployment rate.

[6] All transition rates have been adjusted for margin error. Additionally, the gray vertical bars represent Canadian recessions in my sample.

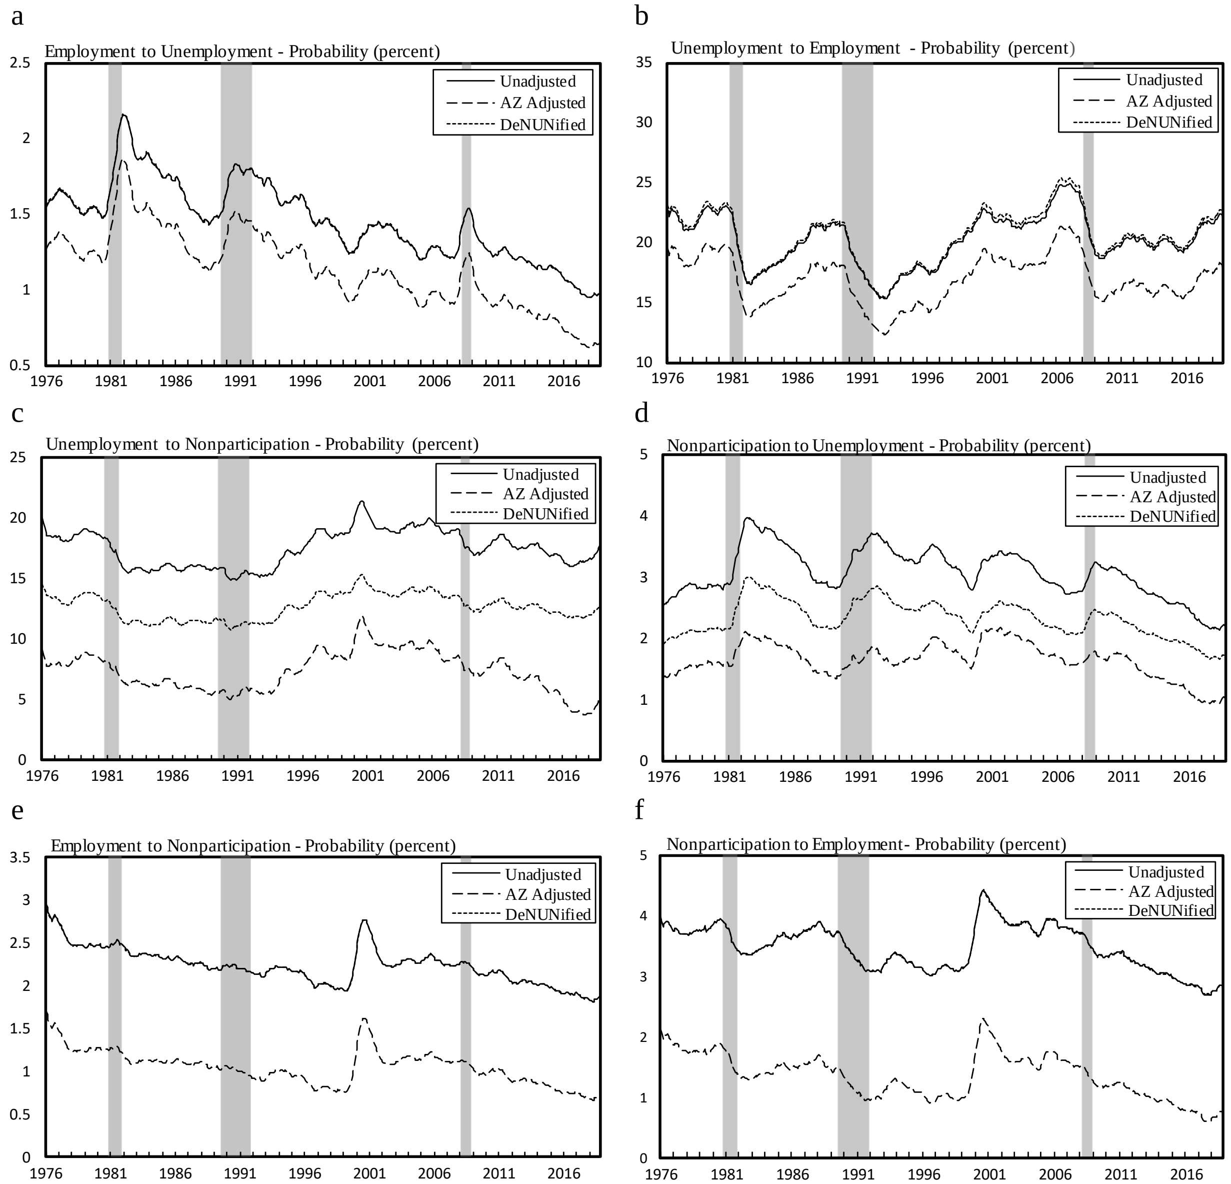

Figure 2. Discrete time monthly transition rates. These rates are calculated using Canadian LFS data. All stock and flow variables are seasonally adjusted, and all transition rates are adjusted for margin error. The grey bars represent recessions in this sample.[7]

4. Methodology

4.1 Adjusting for classification error

[7] In figure 2a, 2e and 2f, the unadjusted data, and the DeNUNified data provides almost identical transition rates over the entire sample. For this reason, the two lines overlap when graphed.

Classification error occurs in survey data when a respondent incorrectly categorises their labour force status. A consequence of this issue is that it can artificially increase the number of observed flows between months. For example, if an individual is a nonparticipant for two consecutive months, but mistakenly is reported to be unemployed in the second month, one would incorrectly observe a N-to-U transition that did not occur. Additionally, regardless of their labour force status in the third month, one would incorrectly report them as transitioning to one of the three states from unemployment, causing a second inaccuracy in the data.

Theoretically, since the distinction between unemployment and nonparticipation can be subtle, they are more susceptible to classification error. A working age individual who is not employed is considered unemployed if they are actively looking for work, or a nonparticipant if they are not actively looking for work. In some cases, the behaviour of an unemployed individual can be very similar to a nonparticipant. Therefore, classification error in survey data is likely to disproportionally overestimate transitions between those two categories, and it is important to address the issue and try to eliminate any bias. In the following sections I will introduce two methods I use to address classification error and explain how they work.

4.1.1 Abowd and Zellner (1985) correction

This first method exploits reinterview survey data to estimate the probability that an individual's true labour force status differs from their initial survey response. To ensure the robustness of a national survey, such as the CPS in the United States, governments will reinterview a subset of the initial sample and analyze inconsistencies. Abowd and Zellner (1985) used this data to estimate the probability of misclassification of individuals across labour-force status.

Denoting (ε_{ij}) as an individual with true labour market status i, but reported as j, you can use the following matrix equation to estimate the true stocks (E, U, N) from their measured stocks $(\widehat{E}, \widehat{U}, \widehat{N})$:

$$\begin{bmatrix} \widehat{E} \\ \widehat{U} \\ \widehat{N} \end{bmatrix}_t = \begin{bmatrix} 1-\varepsilon_{EU}-\varepsilon_{EN} & \varepsilon_{UE} & \varepsilon_{NE} \\ \varepsilon_{EU} & 1-\varepsilon_{UE}-\varepsilon_{UN} & \varepsilon_{NU} \\ \varepsilon_{EN} & \varepsilon_{UN} & 1-\varepsilon_{NE}-\varepsilon_{NU} \end{bmatrix} \begin{bmatrix} E \\ U \\ N \end{bmatrix}_t . \tag{1}$$

By isolating the vector containing the true value of labour market stocks, one gets

$$\begin{bmatrix} E \\ U \\ N \end{bmatrix}_t = \begin{bmatrix} 1-\varepsilon_{EU}-\varepsilon_{EN} & \varepsilon_{UE} & \varepsilon_{NE} \\ \varepsilon_{EU} & 1-\varepsilon_{UE}-\varepsilon_{UN} & \varepsilon_{NU} \\ \varepsilon_{EN} & \varepsilon_{UN} & 1-\varepsilon_{NE}-\varepsilon_{NU} \end{bmatrix}^{-1} \begin{bmatrix} \widehat{E} \\ \widehat{U} \\ \widehat{N} \end{bmatrix}_t , \tag{2}$$

which shows how one can use known observable values of labour market stocks to estimate the underlying true stock values without classification error.

Similarly, Poterba and Summers (1986) showed that the same correction matrix can be used to estimate the true flows (EE, EU, EN, ect ...) from their observed values ($\widehat{EE}, \widehat{EU}, \widehat{EN}, ect$ ...). Using the following three matrices:

$$\boldsymbol{E} = \begin{bmatrix} 1-\varepsilon_{EU}-\varepsilon_{EN} & \varepsilon_{UE} & \varepsilon_{NE} \\ \varepsilon_{EU} & 1-\varepsilon_{UE}-\varepsilon_{UN} & \varepsilon_{NU} \\ \varepsilon_{EN} & \varepsilon_{UN} & 1-\varepsilon_{NE}-\varepsilon_{NU} \end{bmatrix}, \tag{3}$$

$$\boldsymbol{N}_t = \begin{bmatrix} EE & UE & NE \\ EU & UU & NU \\ EN & UN & NN \end{bmatrix}_t, \tag{4}$$

and

$$\widehat{\boldsymbol{N}}_t = \begin{bmatrix} \widehat{EE} & \widehat{UE} & \widehat{NE} \\ \widehat{EU} & \widehat{UU} & \widehat{NU} \\ \widehat{EN} & \widehat{UN} & \widehat{NN} \end{bmatrix}_t, \tag{5}$$

one can begin with an equation outlining how the number of observed transitions between labour market states are based on the true number of transitions, and the rates of classification error:

$$\widehat{\boldsymbol{N}}_t = \boldsymbol{E}\boldsymbol{N}_t\boldsymbol{E}'. \tag{6}$$

By isolating $\boldsymbol{N}_t$ to get,

$$\boldsymbol{N}_t = \boldsymbol{E}^{-1}\widehat{\boldsymbol{N}}_t(\boldsymbol{E}^{-1})', \tag{7}$$

one obtains an equation to estimate the true number of transitions in a labour market using observable data.[8]

One important assumption is that rates of classification error are constant over time. This allows one to use the same correction matrix for all months in their sample. Additionally, this method relies on a strong assumption that reinterviews provide the individuals true status. However, for my paper, I will not assume that the correction matrix for American CPS data will be valid for Canadian LFS data. That is why I will use Canadian LFS reinterview data from Jan 1987 to Nov 1989, provided by table 4 in Singh and Rao (1991), to formulate a correction matrix that can be applied to the Canadian LFS. Table 1.a presents the correction matrix from CPS reinterview data, and table 1.b presents the correction matrix for LFS reinterview data.

[8] Since I am using estimates of prevalence of classification error in my data (ε_{ij}), I am only able to estimate the true stock of workers and the number of transitions in the labour market each year.

Table 1.a
Abowd and Zellner (1985) estimates of classification error (%).

Original interview status	Status determined on reinterview		
	Employed	Unemployed	Non-participant
Employed	98.78	1.91	0.50
Unemployed	0.18	88.57	0.29
Non-Participant	1.03	9.52	99.21

(Source: Abowd and Zellner (1985))

Table 1.b
Estimates of classification error (%), based on table 4 in Singh and Rao (1991) .

Original interview status	Status determined on reinterview		
	Employed	Unemployed	Non-participant
Employed	98.94	1.95	0.69
Unemployed	0.30	90.75	0.74
Non-Participant	0.76	7.30	98.57

(Source: Singh and Rao (1991))

4.1.2 Recoding NUN and UNU sequences.

The second method for addressing classification error is a novel technique developed by Elsby et al. (2015) that can be used in addition to the Abowd-Zellner correction. One advantage of this new method is that it does not rely on the hard to prove assumptions that reinterview data is perfectly accurate, nor that classification error is constant over time. Instead, they make minimal modifications only to the sets of transitions that are the most likely to be inaccurate. If an individual is reported as transitioning from unemployed to nonparticipant and back to unemployed, or from nonparticipation to unemployed and back to nonparticipation, it is likely that their behaviour did not change in a significant way. For this reason, they recode three-month sequences where an individual is unemployed or nonparticipant in both the first and third month but switch to the other state in month two. To apply this technique to LFS data, I need to consider all permutations of six months sequences where an individual is only reported as unemployed once, and it transitions from and to nonparticipation. Likewise, I look for instances of the opposite occurring.

Table 2
Recoding of unemployment-nonparticipation cyclers: “deNUNified” flows.

	Measured	Recoded
UNUs	UNU$x_1x_1x_1$	UUU$x_1x_1x_1$
	x_1UNUx_1x_1	x_1UUUx_1x_1
	x_1x_1UNUx_1	x_1x_1UUUx_1
	$x_1x_1x_1$UNU	$x_1x_1x_1$UUU
NUNs	NUN$x_2x_2x_2$	NNN$x_2x_2x_2$
	x_2NUNx_2x_2	x_2NNNx_2x_2
	x_2x_2NUNx_2	x_2x_2NNNx_2
	$x_2x_2x_2$NUN	$x_2x_2x_2$NNN

Note: The notation ABCDEF refers to a sequence of transitions associated with six consecutive monthly individual labour market states.
x1 represents either E, U, or a missing observation
x2 represents either E, N, or a missing observation
x1 and x2 in each individual case remains unchanged

Table 2 provides a comprehensive list how I recoded six-month sequences depending on the measured data. Elsby et al. (2015)’s deNUNified method should used jointly with other methods to gauge robustness of your results.

4.2Adjustments for temporal aggregation bias.

The second well documented issue with gross flow data is that I only observe labour market status at discrete points one month apart. I cannot observe any changes that occur between survey dates. Let's say individuals are surveyed on the 15th day of each month. A respondent could be employed on January 15st, then transition to unemployment and find a new job all before being surveyed again on February 15st. In this scenario, the raw date would overestimate the number of individuals who do not transition out of employment and underestimate the number of individuals who transition between employment and unemployment. This is just one example of the infinite possible scenarios that could occur when someone is measured as employed in two consecutive months.

Shimer (2012) has provided a solution to this problem which has been used in many labour economics papers. It uses an eigen decomposition to estimate the underlying continuous time flow hazard rates from the observed discrete time monthly transition rates. It works since both rates share the same eigen values but differ in their eigen vectors. Elsby et al. (2015) provide the mathematical process of obtaining the hazard rates in section A3 of their appendix.

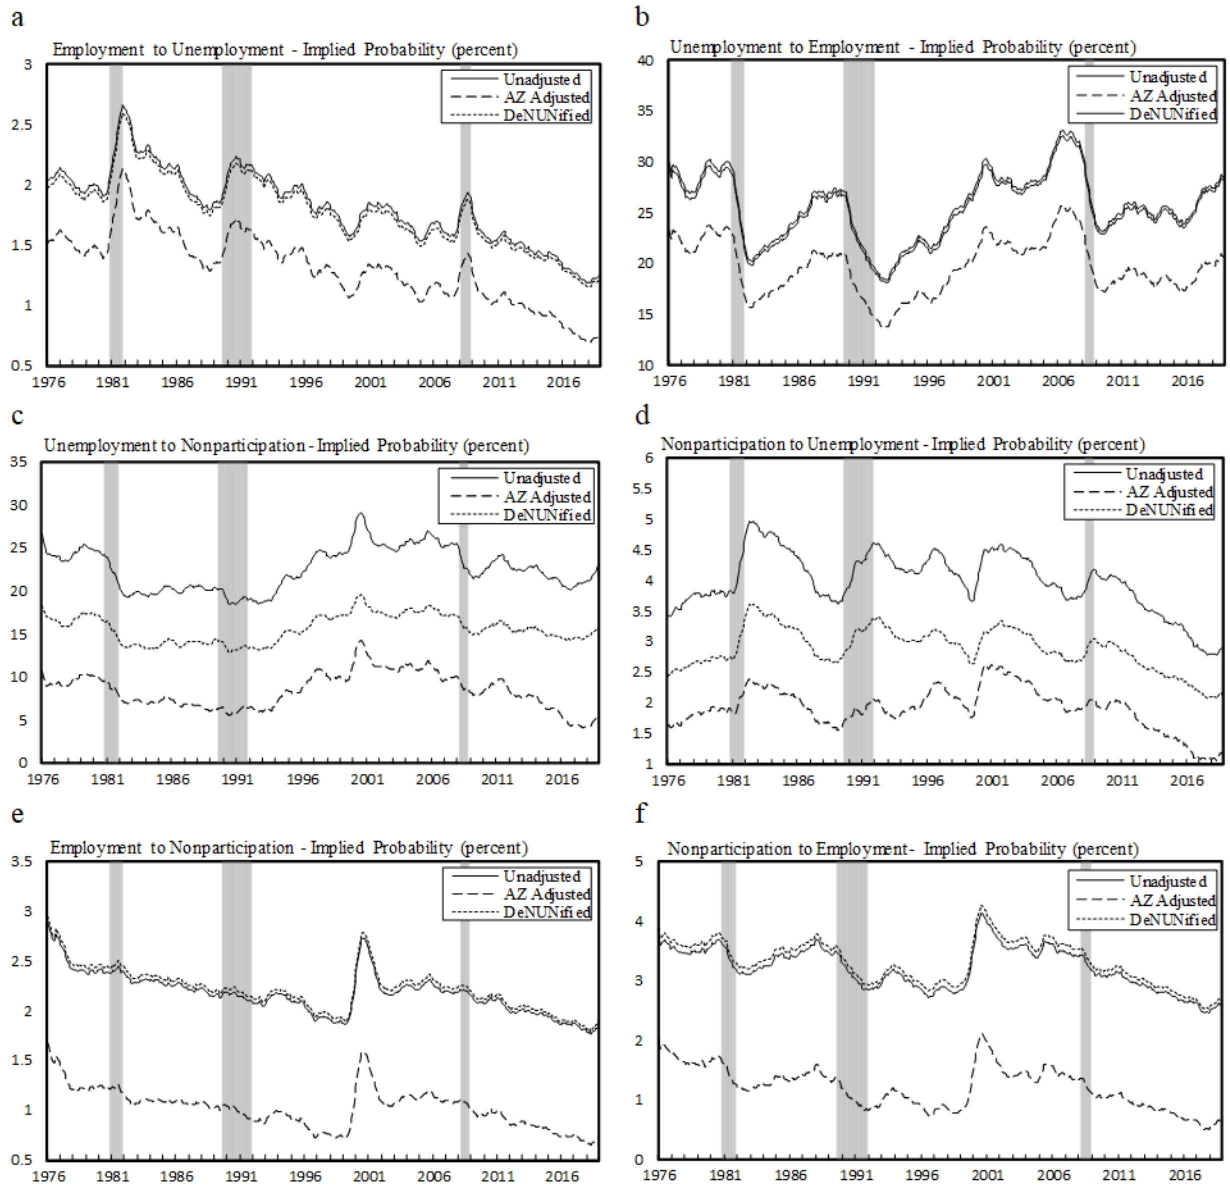

Figure 3. Implied monthly flow hazard rates. These rates are calculated using Canadian LFS data. All stock and flow variables are seasonally adjusted, and all transition rates are adjusted for margin error. The grey bars represent recessions in this sample.

Figure 3 graphs the monthly hazard rates both for the unadjusted data, and the classification error adjusted data. Correcting for both types of errors do not eliminate the cyclical patterns observed in the raw monthly transition rate shown in figure 2. By comparing figure 3 with figure 2, one can see the extend to which temporal aggregation bias causes the raw data to underestimate monthly transition rates. Both the unemployment to employment hazard rate and the unemployment to nonparticipation hazard rate are between five and eight percentage points higher than their associated discrete time monthly transition rate.

The employment to unemployment hazard rate, and the nonparticipation to unemployment hazard rate each are almost one percentage point higher than their discrete time monthly transition rate. It is only transitions between employment and nonparticipation that appear largely unaffected by temporal aggregation bias.

4.3 Margin Error correction.

The final issue with survey data that I will address is margin error. Transition rates are calculated by determining the proportion of employed, unemployed or nonparticipants in the previous month, who transition to any given category in the current month. For any given month, I will not be able to observe the labour market status of every individual in the following month. For that reason, my calculations of the transition rate from employment, or unemployment, or nonparticipation will not sum to 100%. To correct this issue, I will utilize the margin-error correction method developed by Elsby et al. (2013). Using only observed variables, they apply a "weighted-restricted-least-squared adjustment method" to select a vector of margin error adjusted transition probabilities that are consistent with the observed change in the stocks of employed, unemployed and nonparticipating workers for each month.

To derive a vector for estimates of the true monthly transition rates ($\boldsymbol{p}$), one can begin with the equation for the change in labour market stocks:

$$\Delta \boldsymbol{s}_t = \boldsymbol{X}_{t-1}\boldsymbol{p} \tag{8}$$

Which can be expressed as:

$$\begin{bmatrix} \Delta e \\ \Delta u \end{bmatrix}_t = \begin{bmatrix} -e_{t-1} & -e_{t-1} & u_{t-1} & 0 & u_{t-1} & 0 \\ e_{t-1} & 0 & -u_{t-1} & -u_{t-1} & 0 & u_{t-1} \end{bmatrix} \begin{bmatrix} p_{eu} \\ p_{en} \\ p_{ue} \\ p_{un} \\ p_{ne} \\ p_{nu} \end{bmatrix}. \tag{9}$$

For simplicity, from this point on the stock variables are expressed as shares of the population such that $e_t + u_t + n_t = 1$. Using a consistent estimate of the covariance matrix ($\boldsymbol{W}$) for the observed transition rate matrix ($\widehat{\boldsymbol{p}}$), one can set up the weighted-restricted-least-squared minimization problem and derive the associated lagrangian.

$$\boldsymbol{W} = \begin{bmatrix} \frac{\widehat{p_{eu}}(1-\widehat{p_{eu}})}{e_{t-1}} & -\frac{\widehat{p_{eu}}\widehat{p_{en}}}{e_{t-1}} & 0 & 0 & 0 & 0 \\ -\frac{\widehat{p_{eu}}\widehat{p_{en}}}{e_{t-1}} & \frac{\widehat{p_{en}}(1-\widehat{p_{en}})}{e_{t-1}} & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{\widehat{p_{ue}}(1-\widehat{p_{ue}})}{u_{t-1}} & -\frac{\widehat{p_{ue}}\widehat{p_{un}}}{u_{t-1}} & 0 & 0 \\ 0 & 0 & -\frac{\widehat{p_{ue}}\widehat{p_{un}}}{u_{t-1}} & \frac{\widehat{p_{un}}(1-\widehat{p_{un}})}{u_{t-1}} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{\widehat{p_{ne}}(1-\widehat{p_{ne}})}{n_{t-1}} & -\frac{\widehat{p_{ne}}\widehat{p_{nu}}}{n_{t-1}} \\ 0 & 0 & 0 & 0 & -\frac{\widehat{p_{ne}}\widehat{p_{nu}}}{n_{t-1}} & \frac{\widehat{p_{nu}}(1-\widehat{p_{nu}})}{n_{t-1}} \end{bmatrix}^{-1} \tag{10}$$

The solution to the minimization problem can be expressed as followed:

$$\begin{bmatrix} \boldsymbol{p} \\ \boldsymbol{\mu} \end{bmatrix} = \begin{bmatrix} \boldsymbol{W} & \boldsymbol{X}'_{t-1} \\ \boldsymbol{X}_{t-1} & \boldsymbol{0} \end{bmatrix}^{-1} \begin{bmatrix} \boldsymbol{w}\widehat{\boldsymbol{p}} \\ \Delta \boldsymbol{s}_t \end{bmatrix} \tag{11}$$

In the above equation $\boldsymbol{p}$ is a vector of the true monthly transition rates.[9] This equation also provides the two lagrange variables in the 2x1 vector $\boldsymbol{\mu}$.

5. Measuring the role of the participation margin

Now that I have calculated the instantaneous transition rates, they can be used to estimate the proportion of variation in both the unemployment rate and the labour force participation rate that can be attributed to variation in each individual transition rate.

5.1 Decomposition derivation

For my analysis, I will employ a flow-based variance decomposition developed by Elsby et al. (2015). Once again, the stock variables are expressed as shares of the population such that $e_t + u_t + n_t = 1$. This allows me to write the two dimensional discrete-time Markov chain that maps the change in the stock variables over time as,

$$\begin{bmatrix} e \\ u \end{bmatrix}_t = \begin{bmatrix} 1 - p_{eu} - p_{en} - p_{ne} & p_{ue} - p_{ne} \\ p_{eu} - p_{nu} & 1 - p_{ue} - p_{un} - p_{nu} \end{bmatrix}_t \begin{bmatrix} e \\ u \end{bmatrix}_{t-1} + \begin{bmatrix} p_{ne} \\ p_{nu} \end{bmatrix}_t, \tag{12}$$

where $\boldsymbol{s}_t = \begin{bmatrix} e \\ u \end{bmatrix}_t$ represents the stock of workers in time t,

[9] For a more detailed explanation of this margin error correction method, including the lagrangian associated with this weighted-restricted-least-squared minimization problem, see the supplementary appendix for Elsby et al. (2013).

$\widetilde{\boldsymbol{P}}_t = \begin{bmatrix} 1 - p_{eu} - p_{en} - p_{ne} & p_{ue} - p_{ne} \\ p_{eu} - p_{nu} & 1 - p_{ue} - p_{un} - p_{nu} \end{bmatrix}_t$ and $\boldsymbol{q}_t = \begin{bmatrix} p_{ne} \\ p_{nu} \end{bmatrix}_t$ are a matrix and vector of monthly transition rates explaining how the stock variables change between period t-1 and t. Therefore, equation (12) can be written in matrix notation as $\boldsymbol{s}_t = \widetilde{\boldsymbol{P}}_t \boldsymbol{s}_{t-1} + \boldsymbol{q}_t$. In a steady state, this Markov chain becomes,

$$\bar{\boldsymbol{s}}_t = \left(\boldsymbol{I} - \widetilde{\boldsymbol{P}}_t\right)^{-1} \boldsymbol{q}_t.\text{[10]} \tag{13}$$

This equation can be rearranged to get,

$$\bar{\boldsymbol{s}}_t = \widetilde{\boldsymbol{P}}_t \bar{\boldsymbol{s}}_t + \boldsymbol{q}_t. \tag{14}$$

One can subtract equation (14) from the matrix representation of equation (12) to get,

$$\boldsymbol{s}_t - \bar{\boldsymbol{s}}_t = \widetilde{\boldsymbol{P}}_t s_{t-1} - \widetilde{\boldsymbol{P}}_t \bar{\boldsymbol{s}}_t + \boldsymbol{q}_t - \boldsymbol{q}_t, \tag{15}$$

which simplifies to

$$(\boldsymbol{s}_t - \bar{\boldsymbol{s}}_t) = \widetilde{\boldsymbol{P}}_t (\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_t). \tag{16}$$

It will be useful to expand and simply equation (16) in the following way;

$$(\boldsymbol{s}_t - \bar{\boldsymbol{s}}_t) = \widetilde{\boldsymbol{P}}_t (\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_t) + \widetilde{\boldsymbol{P}}_t (\bar{\boldsymbol{s}}_t - \bar{\boldsymbol{s}}_{t-1}) - \widetilde{\boldsymbol{P}}_t \Delta \bar{\boldsymbol{s}}_t, \tag{17}$$

$$(\boldsymbol{s}_t - \bar{\boldsymbol{s}}_t) = \widetilde{\boldsymbol{P}}_t (\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_{t-1}) - \widetilde{\boldsymbol{P}}_t \Delta \bar{\boldsymbol{s}}_t. \tag{18}$$

For the next step, I need to use the following equation of the decomposed change in s_t:

$$\Delta \boldsymbol{s}_t = (\boldsymbol{s}_t - \bar{\boldsymbol{s}}_t) - (\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_{t-1}) + \Delta \bar{\boldsymbol{s}}_t. \tag{19}$$

Substitute $(\boldsymbol{s}_t - \bar{\boldsymbol{s}}_t)$ from (18) into (19) to get,

$$\Delta \boldsymbol{s}_t = \widetilde{\boldsymbol{P}}_t (\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_{t-1}) - \widetilde{\boldsymbol{P}}_t \Delta \bar{\boldsymbol{s}}_t - (\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_{t-1}) + \Delta \bar{\boldsymbol{s}}_t, \tag{20}$$

which simplifies to

$$\Delta \boldsymbol{s}_t = -\left(\boldsymbol{I} - \widetilde{\boldsymbol{P}}_t\right)(\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_{t-1}) + \left(\boldsymbol{I} - \widetilde{\boldsymbol{P}}_t\right) \Delta \bar{\boldsymbol{s}}_t. \tag{21}$$

Pre-multiplying equation (18) by $\widetilde{\boldsymbol{P}}_t^{\,-1}$, one can deduce that $\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_{t-1} - \Delta \bar{\boldsymbol{s}}_t = \widetilde{\boldsymbol{P}}_t^{\,-1} (\boldsymbol{s}_t - \bar{\boldsymbol{s}}_t)$. Therefore, equation (21) can be written as,

$$\Delta \boldsymbol{s}_t = -\left(\boldsymbol{I} - \widetilde{\boldsymbol{P}}_t\right) \widetilde{\boldsymbol{P}}_t^{\,-1} (\boldsymbol{s}_t - \bar{\boldsymbol{s}}_t). \tag{22}$$

Using (22), one can isolate $(\boldsymbol{s}_t - \bar{\boldsymbol{s}}_t)$ and iterate the equation back by one period to get an equation for $(\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_{t-1})$:

$$(\boldsymbol{s}_t - \bar{\boldsymbol{s}}_t) = -\widetilde{\boldsymbol{P}}_t \left(\boldsymbol{I} - \widetilde{\boldsymbol{P}}_t\right)^{-1} \Delta \boldsymbol{s}_t, \tag{23}$$

$$(\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_{t-1}) = -\widetilde{\boldsymbol{P}}_{t-1} \left(\boldsymbol{I} - \widetilde{\boldsymbol{P}}_{t-1}\right)^{-1} \Delta \boldsymbol{s}_{t-1}. \tag{24}$$

Finally, one can replace $(\boldsymbol{s}_{t-1} - \bar{\boldsymbol{s}}_{t-1})$ in equation (21) using equation (24) to get

$$\Delta \boldsymbol{s}_t = \boldsymbol{A}_t \Delta \bar{\boldsymbol{s}}_t + \boldsymbol{B}_t \Delta \boldsymbol{s}_{t-1}, \tag{25}$$

[10] Note that the steady state stock vector $(\bar{\boldsymbol{s}}_t)$ has subscript 't' because changes in the flow hazard rates (f_{ij}) causes shifts in steady state stocks. Therefore, for this decomposition, I want to observe how the steady state stocks change over time.

where $\boldsymbol{A}_t = \left(\boldsymbol{I} - \widetilde{\boldsymbol{P}}_t\right)$ and $\boldsymbol{B}_t = \left(\boldsymbol{I} - \widetilde{\boldsymbol{P}}_t\right)\widetilde{\boldsymbol{P}}_{t-1}\left(\boldsymbol{I} - \widetilde{\boldsymbol{P}}_{t-1}\right)^{-1}$.

Equation (25) above breaks down changes in s_t into two components. $A_t \Delta \bar{s}_t$ represents how labour market stocks are affected by shifts in the steady state cause by changes in the monthly transition rate. $B_t \Delta s_{t-1}$ shows the residual effect of past changes in labour market stocks on its current change.

By iterating this equation backwards, one can express current changes in labour market stock as a function of all past changes:

$$\Delta \boldsymbol{s}_t = \sum_{k=0}^{t-1} \boldsymbol{C}_{k,t} \Delta \bar{\boldsymbol{s}}_{t-k} + \boldsymbol{D}_t \Delta \boldsymbol{s}_0, \tag{26}$$

where $\boldsymbol{C}_{k,t} = (\Pi_{n=0}^{s-1} \boldsymbol{B}_{t-n})\boldsymbol{A}_{t-k}$ and $\boldsymbol{D}_t = (\Pi_{n=0}^{t-1} \boldsymbol{B}_{t-k})$.

As mentioned in footnote six, changes in the flow hazard rates causes shifts in steady state stocks. One can estimate this relationship by utilizing the following approximation for the change in steady state labour market stocks:

$$\Delta \bar{\boldsymbol{s}} \approx \Sigma_{i \neq j} \frac{\partial \bar{\boldsymbol{s}}_t}{\partial f_{ij_t}} \Delta f_{ij_t}, \tag{27}$$

where f_{ij_t} is the continuous time flow hazard rate for transition ij in period t.

The required derivatives in this approximation can be obtained from the continuous-time Markov chain,

$$\dot{\boldsymbol{s}}_t = \begin{bmatrix} -f_{eu} - f_{en} - f_{ne} & f_{ue} - f_{ne} \\ f_{eu} - f_{nu} & -f_{ue} - f_{un} - f_{nu} \end{bmatrix}_t \boldsymbol{s}_t + \begin{bmatrix} f_{ne} \\ f_{nu} \end{bmatrix}_t. \tag{28}$$

Using the equations above, one can create the following flow-based variance decomposition of the unemployment rate:

$$\beta_{ij}^{ur} = \frac{cov\left(\Delta ur_t\, , ur_t(e_t + ur_t)^2 * \ \Delta \boldsymbol{s}_{t,ij}[1,1] + \frac{e_t}{(e_t + ur_t)^2} \Delta \boldsymbol{s}_{t,ij}[2,1]\right)}{var(\Delta ur_t)}, \tag{29}$$

where

$$\Delta \boldsymbol{s}_{t,ij} \approx \left(\boldsymbol{I} - \widetilde{\boldsymbol{P}}_t\right) \sum_{k=0}^{t-1} (\Pi_{n=0}^{s-1} \boldsymbol{B}_{t-n}) \boldsymbol{A}_{t-k} \left(\frac{\partial \bar{\boldsymbol{s}}_{t-k}}{\partial f_{ij_{t-k}}} \cdot \Delta f_{ij_{t-k}}\right) + (\Pi_{n=0}^{t-1} \boldsymbol{B}_{t-k}) \Delta \boldsymbol{s}_0,$$

and e_t and ur_t respectively are the employment-to-population ratio and unemployment rate.

5.2 Results - Replication

To ensure that my code does not have any flaws, I first replicate table 3 from Elsby et al. (2015), before applying my code to the Canadian labour market.[11] The table 3 shows the results I have produced

[11] Elsby et al. (2015) – table 3, page 74.

using the author's original data.[12] My code was able to reproduce the variance decomposition for both the unemployment rate and the labour force participation rate within five or more digits after the decimal. From this replication, I can now apply this decomposition to the Canadian labour market.

[12] Rows 5 and 11 in table 3 could not be calculated since Elsby et al. (2015) only had DeNUNified data starting in 1978.

Table 3 - Replication of table 3 in Elsby et al. (2013) using my code with their data

Variance decomposition of the Canadian Unemployment Rate based on classification error

Type of Classification error adjustment	Sample Period	Change Decomposed	Share of Variance							Total variance between		
			EU	*UE*	*UN*	*NU*	*EN*	*NE*	residual	*E* and *U*	*N* and *U*	*E* and N
Unadjusted	1967 - 2012	Monthly Change	21.3	34.8	22.5	17.4	-0.2	0.6	3.6	56.1	0.4	39.9
DeNUNified	1967 - 2012	Monthly Change	30.5	33.6	16.0	15.9	-0.3	0.5	3.9	64.1	0.2	31.9
Abowd-Zellner	1967 - 2012	Monthly Change	26.2	38.7	22.6	11.9	-0.3	0.5	0.5	64.8	0.2	34.5
Unadjusted	1967 - 2012	Quarterly Average	24.9	34.9	23.9	9.5	-0.3	1.0	6.0	59.8	0.8	33.4
DeNUNified	1967 - 2012	Quarterly Average	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Abowd-Zellner	1967 - 2012	Quarterly Average	29.6	41.7	26.7	-0.7	-1.3	2.1	1.8	71.4	0.8	26.0
Unadjusted	1978 - 2012	Quarterly Average	22.3	35.1	22.3	13.2	-0.7	1.5	6.3	57.4	0.8	35.5
DeNUNified	1978 - 2012	Quarterly Average	25.2	42.5	17.1	11.6	-0.8	1.1	3.2	67.7	0.3	28.7
Abowd-Zellner	1978 - 2012	Quarterly Average	25.6	44.4	26.4	3.9	-1.7	2.3	-0.9	70.0	0.7	30.3
Unadjusted	1967 - 2012	12-Month Change	19.1	38.4	23.7	10.7	-0.2	1.1	7.1	57.6	0.9	34.4
DeNUNified	1967 - 2012	12-Month Change	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Abowd-Zellner	1967 - 2012	12-Month Change	23.7	46.3	29.0	-0.9	-0.7	1.1	1.4	70.0	0.4	28.1
Unadjusted	1978 - 2012	12-Month Change	17.2	39.3	22.3	12.5	-0.7	1.6	7.7	56.5	0.9	34.9
DeNUNified	1978 - 2012	12-Month Change	20.5	47.9	18.9	9.9	-0.7	1.1	2.4	68.4	0.4	28.8
Abowd-Zellner	1978 - 2012	12-Month Change	19.7	49.6	28.2	2.5	-1.4	1.6	-0.4	69.3	0.3	30.8

Note: Elsby et al. (2015) have provided an excel file with United States CPS data and their own code they use to produce table 3 in their paper. I have written my own code in Rstudio to run this decomposition. This table shows the results of my code for the variance decomposition outlined in section 5.1, using Elsby et al. (2015)'s CPS data on the United States labour market. Stocks and flows have been seasonally adjusted. All monthly transition rates, and hazard rates have been adjusted for margin error correction.

Table 4

Variance decomposition of the Canadian Unemployment Rate based on classification error

Type of Classification error adjustment	Sample Period	Share of Variance							Total variance between		
		EU	*UE*	*UN*	*NU*	*EN*	*NE*	residual	*E* and *U*	*N* and *U*	*E* and N
Unadjusted	1976 - 2019	20.73	34.40	13.91	21.06	0.74	0.87	8.28	55.13	34.98	1.61
DeNUNified	1976 - 2019	23.42	38.08	10.99	18.21	0.58	0.58	8.14	61.50	29.20	1.16
Abowd-Zellner	1976 - 2019	29.34	38.63	9.26	14.36	0.73	-0.09	7.76	67.97	23.62	0.65
Unadjusted	1978 - 2012	20.71	34.39	13.85	21.14	0.74	0.88	8.29	55.10	34.99	1.62
DeNUNified	1978 - 2012	23.40	38.07	10.96	18.27	0.58	0.59	8.14	61.47	29.22	1.17
Abowd-Zellner	1978 - 2012	29.30	38.62	9.19	14.48	0.73	-0.08	7.77	67.91	23.67	0.65

Note: For this decomposition I use Canadian gross flow data from the Labour Force Survey between January 1976 and December 2019. All stock and flow variables have been seasonally adjusted and all hazard rates have been obatained using margin adjusted-monthly transition rates. Additionally, I use the quarterly average of the unemployment rate, and flow hazard rates for this decompotion.

5.3 Results – Canada

Table 4 shows the results of my decomposition of the quarterly average of the unemployment rate for the six flow hazards (f_{ij}) derived in section 4.2. Rows one to three covers the full sample, while rows four to six covers the restricted sample used in Elsby et al. (2015).[13] Table 4 includes the results obtained when using unadjusted flows, and when adjusting for classification error.

Looking at row 1, you can see that countercyclical job loss, and procyclical job gain respectively explain 21 and 34 percent of the change in unemployment rate between 1976 and 2019. Together, all transitions between employment and unemployment can explain a majority of the fluctuations in the unemployment rate. Relative to the United States, E-to-U transitions play less of a role influencing the unemployment rate, and U-to-E transitions play an increased role.

Looking at columns 3 and 4, you can see that labour force participation also plays a significant role in explaining unemployment rate cycles. Leaving the labour force from unemployment and entering unemployment from nonparticipation respectively account for around 14 and 21 percent of unemployment rate fluctuations, combining for a total contribution of roughly 35%. Although the share of total variance between those two transition rates is similar when using American data, Elsby et al. (2013) find that U-to-N transitions play a larger role than N-to-U transitions. Using Canadian data, I find the opposite relationship.

Additionally, rows 5 and 6 shows that transitions between employment and nonparticipation play an insignificant role in accounting for fluctuations in the unemployment rate.

To determine the robustness of these results, I can compare the results from the unadjusted data to rows 2 and 3, which presents the decomposition results using classification error adjusted hazard rates based on Abowd and Zellner (1985)'s method and using DeNUNified flows developed by Elsby et al. (2015). Relative to the unadjusted hazard rates, the share of unemployment variation at can be attributed to transitions from employment to unemployment, and unemployment to employment together increases by 6 percentage points when using DeNUNified flows, and 13 percentage points when using AZ adjusted flows. Additionally, I find that transitions from nonparticipation to unemployment, and unemployment to nonparticipation collectively decreases by 6 percentage points when using DeNUNified flows, and 11 percentage points when using AZ adjusted flows. The impact of flows between nonparticipation and employment decreases slightly and remain insignificant in explaining changes in unemployment.

This result is expected since classification error primarily affects classification in unemployment and nonparticipation, and, as a result, survey data tends to overestimate transitions between those two

---

[13] The full sample includes data from 1976 to 2019, while the restricted sample only includes data from 1978 to 2012. The restricted sample allows me to directly compare my results using Canadian data, with the results of Elsby et al. (2015), who use US data. Both of our papers include a sample from 1978 to 2012 in our decomposition.

categories. However, I find that the participation margin still plays a significant role accounting for 35 to 24 percent of the variation in the Canadian unemployment rate.

It is worthwhile to compare my results with past papers that used a two-state framework to decompose the unemployment rate on the job finding and job separation rate in Canada (Campolieti (2011), Elsby et al. (2013)). Both papers used unemployment duration data to calculate the monthly outflow rate using a method developed by Shimer (2012) and find an 80-20 outflow-inflow split in accounting for unemployment fluctuations. However, my findings indicate that unemployment inflows (from nonparticipation and employment) play a significantly larger role. Using $f_{UE} + f_{UN}$ as a measure of unemployment outflows in a three-state model, and $f_{EU} + f_{NU}$ as a measure of unemployment inflows, I find that outflows and inflows are responsible for around 48% and 42% respectively.[14]

One reason for this difference is that I use gross flow data instead of unemployment duration data. This discrepancy is also observed in papers analyzing the American labour market. By comparing Elsby et al. (2013) with Elsby et al. (2015), one can notice a similar result. The former finds an 85-15 outflow-inflow split using duration data, while the later finds a two thirds – one thirds split using gross flow data.[15] Although this is the first paper to use Canadian gross flow data in an unemployment rate variance decomposition, and finds significantly different results than Campolieti (2011) or Elsby et al. (2013), my results are not unexpected since they follow the same pattern as seen between American unemployment variance decompositions when using gross flow, and unemployment duration data.[16]

6. Dynamics of the Canadian Economy

When comparing the results of this decomposition using Canadian data to the same decomposition using American data, one can notice that the residual for the Canadian decomposition is consistently higher. This is evidence that the Canadian labour market is less dynamic than its American counterpart.[17] Since this decomposition method measures the covariance between the observed unemployment rate, and the contribution of each flow hazard rate to changes in steady state unemployment, the size of the residual is mainly a reflection of the difference between the true unemployment rate, and its steady state. Therefore,

[14] Additionally, around 2% is explained by $f_{EN} + f_{NE}$, and the remaining 8% is unexplained in the residual.
[15] The difference in decomposition results when using gross flow data, and unemployment duration data is certainly worth investigating, however it is beyond the scope of this paper.
[16] Additionally, the difference between my results and past papers could partially be caused by my use of a three-state model which accounts for labour force growth and distinguish between inflows from employment or nonparticipation. However, as mentioned in Shimer (2012) and Elsby et al. (2013), this will not change my results in a significant way.
[17] The extent to which an economy is dynamic refers to their elasticity in response to shocks. A more dynamic economy is more elastic, and will efficiently make adjustments to move towards their new steady state.

my results indicate that the Canadian labour market is slower to adjust when experiencing shocks and on average, will have a larger difference between the observed unemployment rate and its steady state.

This observation supported by Elsby et al. (2013). Using data from 14 OECD economies, they implement two similar variance decompositions to estimate the proportion of variance in the unemployment rate that can be explained by a countries deviation from its steady state. The first decomposition, formulated by Fujita and Ramey (2009), only accounts for variation in the job finding rate, and the job separation rate. Elsby, Hobijn and Şahin modify that decomposition to include the country's initial deviation from its steady state in the first time period as a third source of variation. If the residuals of the two decompositions are very close, then the economy is very dynamic and adjust to shocks quickly and efficiently moves towards their new steady state. Their data finds that Canada's residual reduced from -5% to -2% in the second decomposition, meanwhile the residual for the United States remains constant in both decompositions at -1%.

Therefore, my findings are consistent with the existing literature by showing that the Canadian Economy is slightly less dynamic than the United States. For future research, modifying the decomposition used in this paper to account for deviations from the steady state would allow it to more accurately explain unemployment rate fluctuations in less dynamic economies.

7. Stock-based decomposition

One takeaway from the flow-based variance decomposition shown in section 5 is that the participation margin, specifically transitions between unemployment and nonparticipation, play a significant role in determining the unemployment rate. However, as done in Elsby et al. (2015), I can also use a simpler stock-based decomposition to determine the portion of changes in the unemployment rate that result from changes in the labour force participation rate and the employment-to-population rate. Their decomposition can be derived from the following approximation of the unemployment rate:

$$\Delta ur_t \approx (1 - ur_{t-1})(\Delta logL_t - \Delta logE_t). \tag{30}$$

From this approximation, they derived the following two beta equations:

$$\beta_L^{ur} = \frac{cov(\Delta ur_t,(1-ur_{t-1})\Delta logL_t)}{var(\Delta ur_t)}, \tag{31}$$

$$\beta_E^{ur} = \frac{cov(\Delta ur_t,(ur_{t-1}-1)\Delta logE_t)}{var(\Delta ur_t)}. \tag{32}$$

In table 5, I show the results of applying this decomposition to my full sample, and to each recession individually. When applied to the full sample, I get a similar result to Elsby et al. (2015). The share of variance attributed to the participation margin (-20%), is significantly lower, and of the opposite sign than what I find in my flow-based decomposition (35%).[18]

Table 5
Stock-based decomposition of the Unemployment Rate on labour force participation rate, and the employment to population ratio

Sample Period	Share of variance (%)		
	Log change in L	Log change in E	residual
Full sample			
Aug 1976 to Jul 2019	-20.12	120.21	-0.09
Recessions			
Jun 1981 to Oct 1982	-33.17	133.44	-0.27
Mar 1990 to May 1992	27.70	72.46	-0.16
Oct 2008 to May 2009	106.69	-6.49	-0.20

Note: This table is calculated using LFS confidential files. All stock variables have been seasonally adjusted.

It is important to understand that changes in the participation margin results from four separate flows, namely, transitions from employment or unemployment to nonparticipation, and vice versa.

The stock-based decomposition is most accurate when most worker flows are not offset by simultaneous flows in the opposite direction. When this is true, the net changes in the stock of workers provide sufficient information to understand why the unemployment rate changes. However, this contradicts the large body of evidence showing that a substantial percentage of worker flows do offset each other. For example, on average between February 1976 and October 1991, 190,000 workers transitioned from employment to unemployment per month and 235,000 transitions from unemployment to employment per month (Jones (1991)). Likewise, transitions between nonparticipation and employment or unemployment have a similar pattern of mostly offsetting flows. Overall, between those months only 7.66% of all flows were not offset in the sample (Jones (1991)).[19] Since changes the participation rate is a byproduct of mostly offsetting flows, the stock-based decomposition is misleading regarding the sources of unemployment fluctuations.

[18] By comparison, Elsby et al. (2015) finds the share of variance attributed to the log participation margin in their stock-based decomposition is around -7%. This provides evidence that the participation rate in slightly more cyclical with unemployment in Canada relative to the United States.
[19] Elsby et al. (2015) shows that this phenomenon is observed in the US labour market, and I provide evidence that it occurs in the Canadian labour market as well.

To further investigate how changes in stocks influence the unemployment rate, I analyze the three Canadian recessions observed in my sample. Table 6 provides the cumulative change in the unemployment rate, the log participation rate, and the log employment-to-population ratio during each recession.

Table 6
Cumulative change in the unemloyment rate, and its components, during the three canadian recessions in our sample

Sample	Cumulative change in		
	Unemployment rate	log L	log E
Jun 1981-Oct 1982	0.0461	-0.0092	-0.0602
Mar 1990-May 1992	0.0336	-0.0225	-0.0595
Oct 2008-May 2009	0.0134	-0.0031	-0.0176

Note: This table is calculated using LFS confidential files. All stock variables have been seasonally adjusted.

During all three recessions, I observe a decrease in both the participation rate, and the employment to population ratio. However, during the 1990/92 recession, I notice that the change in the participation margin, relative to the change in employment, is significantly higher than what is observed during the other two recessions. To understand why this occurs, I will take the next step and analyse worker flows in more details.

Figure 4 a, b and c, shows the estimated contribution of changes in the flow hazard rate towards the unemployment rate for 60 months after the start of each recession. Across all three recessions, figure 4 shows that transitions between unemployment and nonparticipation consistently explain around 30% of the variation in the unemployment rate. So how can one explain table 6's findings that the participation rate plays a comparatively larger role in the 1990/92 recession?

To answer this question, one can examine graphs d, e and f in figure 4. They show the estimated contribution of changes in the flow hazard rate towards the participation rate for 60 months after the start of each recession. Much like in the United States, one observes that changes in Canadian flow hazard rates between employment and unemployment are largely canceled out by opposing changes in hazard rates between unemployment and nonparticipation. The main source of variation in the participation margin results from changes in the hazard rates between employment and nonparticipation. Figure 3 reinforces this finding, since during the 1990/92 recessions, transitions between nonparticipation and employment fell more than either of the other two recessions.

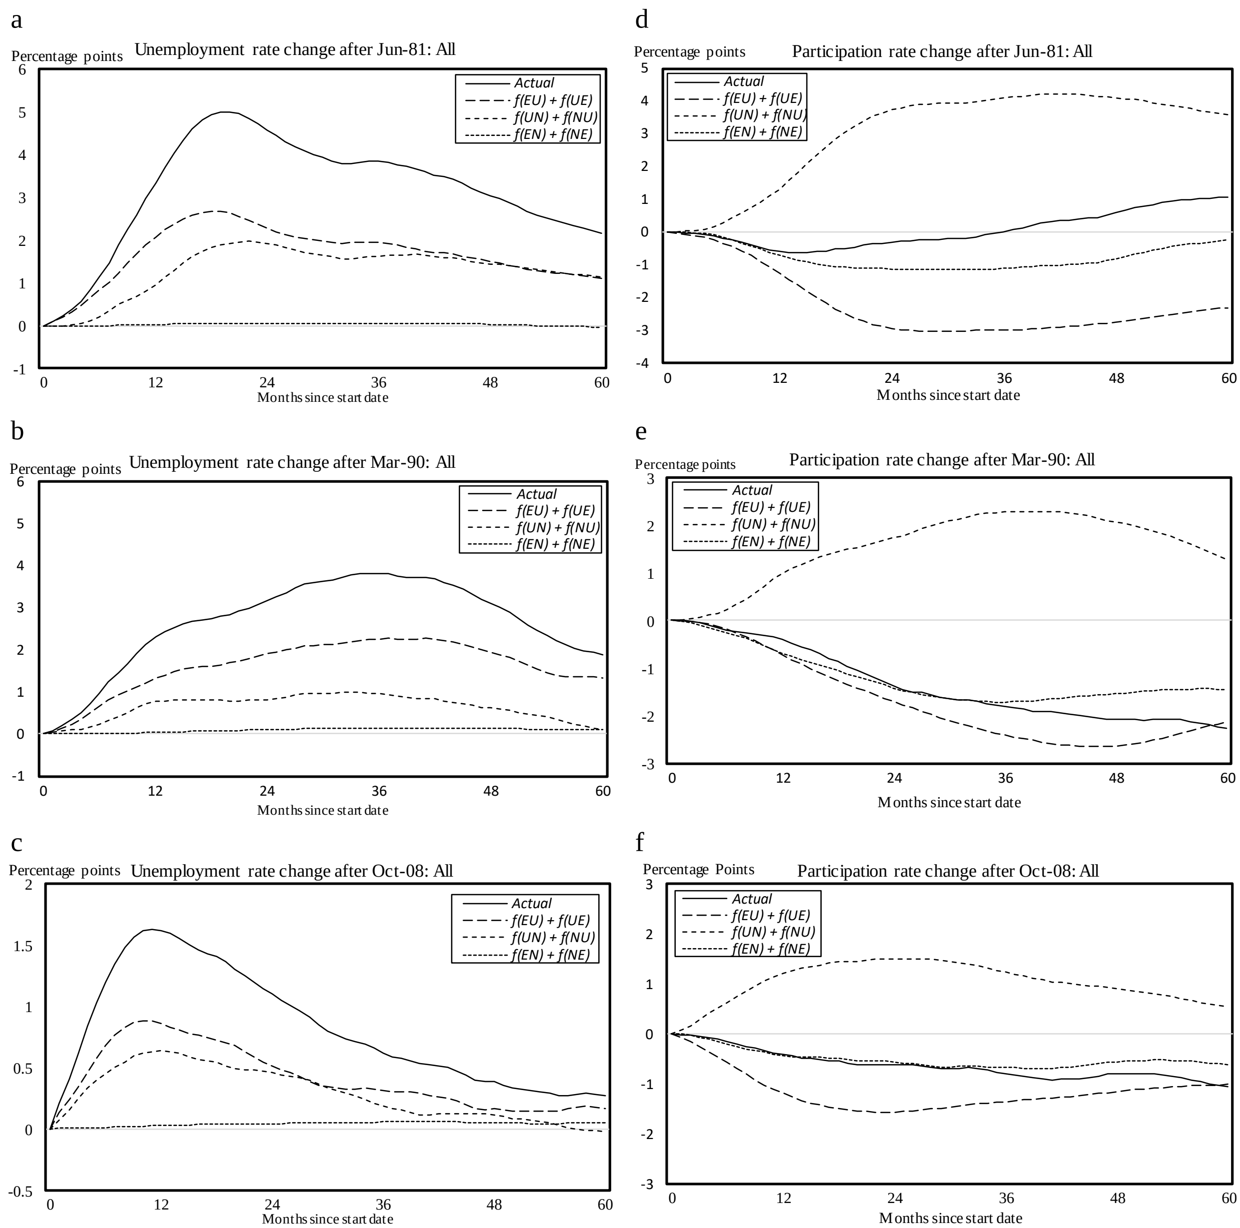

Figure 4. Contributions of flows to changes in the unemployment and labour force participation during recessions. These figures are calculated using Canadian LFS data. All stock and flow variables are seasonally adjusted, and all transition rates are adjusted for margin error.

Therefore, from my analysis of the Canadian data, I can support the statement made in Elsby et al. (2015) that a stock-based decomposition does not accurately explain the impact of the participation margin on changes in unemployment during recessions. Additionally, I show through figure 4 that changes in the participation margins during recessions are mainly a result of changes in the margin between employment to nonparticipation flows and nonparticipation to employment flows. Transitions between E and U; and U and N tend to cancel each other out during recessionary periods.

8. Towards understanding the participation margin, and labour force exit rates

In previous sections, I have used a variance decomposition along with a visual analysis of figure 2 and figure 3 to determine that the participation margin (transitions between unemployment and nonparticipation) does play a significant role in determining unemployment cycles. This is especially true during recessions where the unemployment to nonparticipation hazard rate rise sharply, and gradually decrease afterwards. This phenomenon is robust to corrections for classification error and time aggregation bias.

In the following sections, I try to answer why the participation margin impacts the unemployment rate by analyzing changes in labour market demographics over time.

8.1 Heterogeneity in labour force exit rates.

In this section, I explore how labour force attachment varies across different demographics within the labour market. It is especially useful to identify how labour force attachment affects behavior during recessions.

To better understand how labour force attachment affects the participation margin, I can break down this sample by demographic. In table 7, I present the average monthly U-to-N and U-to-E transition probabilities for the sample, separated by sex, age categories, education level, and reason for unemployment.[20]

[20] In classifying the reason for unemployment, the Canadian LFS includes a special category for workers who have been offered a job that will not start until a future date. I have excluded this category from this sample due to low sample sizes.

Table 7
Heterogeneity in Unemployment exit probability

Subgroup of the Unemployed	Unemployment exit probability (%)			Changes in unemployment share during recession (%)
	P_{UN}	P_{UE}	$P_{UN} + P_{UE}$	
Sex				
Male	15.40	21.38	36.77	3.40
Female	21.02	19.31	40.33	-3.40
Age				
15-24	23.64	23.14	46.78	-1.65
25-54	14.69	20.10	34.80	1.70
55 and over	19.87	14.88	34.75	-0.06
Reason for Unemployment				
Job Loser	12.31	16.66	28.97	5.43
Job Leaver	8.82	24.97	33.79	-2.85
Labour Force Entrant	29.70	17.54	47.24	-2.58

Note: All samples cover the years 1976-2019

Additionally, the final column of table 7 shows the average change in the share of the unemployment pool for each group during a recession. By analyzing the final column of table 7, you can see that demographics which traditionally have lower labour force attachment (females, younger individuals, older individuals, job leavers, labour force entrants, less educated individuals) leave the labour force at higher rates during recessions.

I also observe during recessions that a larger share of the unemployment pool is comprised of individuals who have lost their job. The intuition behind this finding is that the temporary decline in economic conditions will cause employers to minimize expense, partially by reducing their workforce. From the perspective of the employee, since it would be harder to get a job while unemployed, they are more likely to only leave their job voluntarily if they plan on transitioning to nonparticipation. For individuals who want to look for a new job during a recession, but are not let go from their current employer, it is plausible that they would wait until the recession is over, as that is when more opportunities would arise.

Table 8
Heterogeneiety in Unemployment exit probability, separated by recessionary periods and non-recessionary periods

Subgroup of the Unemployed	Unemployment exit probability (%) during recessions			Unemployment exit probability (%) during non-recessions		
	P_{UN}	P_{UE}	$P_{UN} + P_{UE}$	P_{UN}	P_{UE}	$P_{UN} + P_{UE}$
Full Sample	16.71	19.36	36.07	17.98	20.61	38.60
Sex						
Male	13.93	20.19	34.12	15.56	21.51	37.07
Female	20.45	18.26	38.71	21.09	19.43	40.52
Age						
15-24	21.34	21.27	42.61	23.90	23.35	47.25
25-54	14.00	18.92	32.92	14.77	20.23	35.01
55 and over	19.46	15.03	34.48	19.92	14.86	34.78
Reason for Unemployment						
Job Loser	12.37	15.85	28.22	12.30	16.75	29.05
Job Leaver	9.24	22.92	32.16	8.78	25.20	33.97
Labour Force Entrant	28.67	16.27	44.94	29.82	17.68	47.50

Note: The recession sample includes: Jun 1981 to Oct 1982; March 1990 to May 1992; and Oct 2008 to May 2009. The non-recession sample include all other months between August 1976 to June 2019.
As our method for margin error adjustments requires all 9 monthly transition rates, we were not able to correct these transition rates for margin error. This sample is limited to only the unemployed population.

To reinforce this hypothesis, you can look at table 8, which show monthly exit probabilities for each demographic during recessions, and during non-recessionary periods. It shows that job leavers are the only demographic who have a significant increase in monthly unemployment to nonparticipation transition rate during recessions.[21] This reinforces the idea that job leavers have a very low labour force attachment. In response to a decreased probability of finding employment, they leave the labour force at an increased rate.

By compare my results to table 7 in Elsby et al. (2015), you can see that the demographics of the Canadian Labour market change considerably less during a recession than in the United States. In Canada, as a percentage of the unemployed population, males increase by 3.4% compared to 4.7%, individuals ages 25-54 increases 1.7% compared to 5.4%, and job losers increase 5.4% compared to 12.4%. This once again shows that the Canadian labour market is less dynamic.

[21] Job losers also see their unemployment to nonparticipation rate increase, but only by 0.07 percentage points. This is not a significant increase.

8.2 Labour Force Entry

In this section I am going to analyze patterns in labour force entry into unemployment. Table 4 shows that labour force entry into unemployment plays a larger role than labour force exit from unemployment. Similarly, a visual examination of both figure 2 and figure 3 shows that labour force entry is comparatively more volatile. Since labour force entry rates tend to be small (unlike U-to-N transition rates) should ensure that this observation is not a result of any source of bias.

8.2.1 Classification error and time aggregation

It is important to consider that the countercyclicality of the nonparticipation to unemployment monthly transition rates could be caused by either classification error or time aggregation bias.

Classification error is likely to inflate reports of UN and NU transitions, especially during recessions when a larger share of the population falls into these two categories. However, figure 3 shows that adjusting for classification error with either method I used did not eliminate the volatility of these flows, including during recessionary periods. Therefore, I can be confident that classification error cannot fully account for the countercyclicality of N to U transitions.

Similarly, since I only have monthly survey data, it will miss instances of individuals who find jobs between months. It could be the case that new entrants during a recession will have a harder time finding new jobs leading to an increase in the observed N-to-U rate. However, by comparing figures 2 and 3, one sees that controlling for time aggregation bias does not eliminate the cyclical patterns of the N-to-U hazard rate.

Furthermore, by analysing the final column of table 8 to see if the decrease in likelihood for labour force entrants to find employment during recessions is disproportionate to the decrease experienced by the entire population. I find that the monthly unemployment to employment transition probability does decrease by almost one and a half percentage points in during recessions, however, this is similar to the decrease experienced by the entire unemployed population.

8.2.2 Other explanations

This provides context for the variance decomposition in section 5.3. In Canada, flows into unemployment from nonparticipation play a significantly larger role in affecting the unemployment rate, relative to in the United States. This finding is robust to corrections for classification error and temporal aggregation bias. One possible explanation for this finding could be an increased role of the added worker effect. It could be the case that labour force nonparticipants in Canada have higher labour force attachment and are more willing to enter the labour force during recessions. A second possibility is that since a higher percentage of Canadian workers are in the public sector and/or are unionized (Morissette (2022), Bureau of Labor Statistics (2023), Lammam, Palacios, Ren and Clemens (2015)), they have higher job security, and employment varies less during recessions. That could mean that the added worker effect in Canada is not larger than in the US, but simply plays a comparatively larger role due to decreased variation in employment exits. However, answering this question is beyond the scope of this paper and should be the subject of future research.

9. Conclusion

In this paper I show that the participation margin plays a significant role in determining the unemployment rate in Canada. Specifically, countercyclical flows into unemployment from nonparticipation rise sharply during recessions leading to high unemployment. My finding shows the importance of analyzing labour markets through a three-state model in any future research. Additionally, this paper is the first to apply a flow-based variance decomposition of the unemployment rate in Canada that uses gross flow data. By comparing my results to decompositions using unemployment duration data, I find a much more balanced 48-42 outflow-inflow split in accounting for unemployment fluctuations compared to their finding of an 80-20 split. However, this pattern mimics the discrepancy between gross flow and unemployment duration data in studies of the United States labour market. Finally, my findings are consistent with existing literature showing that Canada has a much less dynamic labour market than in the US. As a result, the demographic composition of the unemployed population does not change during recessions to the same extent as in the United States.

References

Abowd, J.M., A. Zellner. (1985) "Estimating Gross Labor-Force Flows." *Journal of Business & Economic Statistics* 3(3): 253-283

Brochu, P. (2021) "A Researcher's guide to the Labour Force Survey: Its Evolution and the Choice of Public Use versus Master Files." *Canadian Public Policy* 47(3): 335-357

Brochu, P., J. Créchet. (2022) "Survey Non-Response in COVID-19 Times: The Case of the Labour Force Survey." *Canadian Public Policy* 48(3): 451–472

Bureau of Labor Statistics. (2023) "Union Members – 2023." US Department of Labor, 1–13

Campolieti, M. (2011) "The Ins and Outs of Unemployment in Canada, 1976–2008." *The Canadian Journal of Economics* 44(4): 1331-1349

Diamond, P.A. (1982) "Aggregate Demand Management in Search Equilibrium." *Journal of Political Economy* 90(5): 881–894

Elsby, M.W.L., B. Hobijn, A. Şahin. (2013) "Unemployment Dynamics in the OECD." *The Review of Economics and Statistics* 95(2): 530-548

Elsby, M.W.L., B. Hobijn, A. Şahin. (2015) "On the Importance of the Participation Margin for Labor Market Fluctuations." *Journal of Monetary Economics* 72: 64–82

Elsby, M.W.L., R. Michaels., G. Solon. (2009) "The Ins and Outs of Cyclical Unemployment." *American Economic Journal: Macroeconomics* 1(1) 84-110

Fienberg, S.E., E.A. Stasny. (1993) "Estimating Monthly Gross Flows in Labour Force Participation." *Survey Methodology* 9(1): 77-102

Fujita, S., G. Ramey (2009) "The Cyclicality of Separation and Job Finding Rates." *International Economic Review* 50(2) 415–430

Jones, S.R.G. (1993) "Cyclical and Seasonal Properties of Canadian Gross Flows of Labour." *Canadian Public Policy* 19(1): 1–17

Jones, S.R.G., C. Riddell. (1999) "The Measurement of Unemployment an Empirical Approach." *Econometrica* 67(1): 147-162

Jones, S.R.G., C. Riddell. (2006) "Unemployment and Nonemployment: Heterogeneities in Labor Market States." *The Review of Economics and Statistics* 88(2) 314-323

Jones, S.R.G., C. Riddell. (2019) "Unemployment, Marginal Attachment and Labor Force Participation in Canada and the United States." *Journal of Labour Economics* 37(S2) S399-S440

Krusell, P., T. Mukoyama., R. Rogerson., A. Şahin. (2011) "A Three State Model of Worker Flows in General Equilibrium." *Journal of Economic Theory* 146: 1107–1133

Krusell, P., T. Mukoyama., R. Rogerson., A. Şahin. (2017) "Gross Worker Flows over the Business Cycle." *American Economic Review* 107(11): 3447–3476

Krusell, P., T. Mukoyama., R. Rogerson., A. Şahin. (2020) "Gross Worker Flows and Fluctuations in the Aggregate Labor Market." *Review of Economic Dynamics* 37: S205-S226

Lammam, C., M. Palacios., F. Ren., J. Clemens. (2015) “Comparing Government and Private Sector Compensation in Canada.” Fraser Institute

Morissette, R. (2022) “Unionization in Canada, 1981 to 2022.”, *Economics and Social Reports* 2(11)

Mortensen, D.T., C. Pissarides. (1994) “Job Creation and Job Destruction in the Theory of Unemployment.” *The Review of Economic Studies* 61(3): 397–415

Pissarides, C.A. (1985) “Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages.” *The American Economic Review* 75(4): 676–690

Poterba, J.M., L.H. Summers. (1986) “Reporting Errors and Labor Market Dynamics.” *Econometrica* 54(6): 1319–1338

Rogeron, R., R. Shimer (2011) “Handbook of Labor Economics - Chapter 7.” 619-700

Shimer, R. (2012) “Reassessing the Ins and Outs of Unemployment.” *Review of Economic Dynamics* 15: 127–148

Singh, A.C., J.N.K. Rao. (1991) “Classification Error Adjustments for Gross Flow Estimates.” Statistics Canada