

Long Time Out: Unemployment and Joblessness in Canada and the United States

Kory Kroft, Fabian Lange, Matthew J. Notowidigdo, and Matthew Tudball*

June 30, 2017

Abstract

We compare patterns of unemployment and joblessness between the U.S. and Canada during the Great Recession. We document a rise in long-term unemployment in Canada and we estimate how the incidence of long-term unemployment varies across demographic groups in Canada. Similar to previous findings for the U.S., we find that changes in the composition of unemployed cannot account for the recent rise in long-term unemployment. We then extend the matching model in Kroft et al. [2016] to exploit the restricted-access panel data from the Canadian Labor Force Survey which contains information on the time since the last job (“joblessness duration”) for both unemployed individuals and non-participants. This allows us to model duration dependence in all labor force flows involving either unemployment or non-participation. To calibrate the extended matching model, we create a new historical vacancy series for Canada based on relative employment in “recruiting industries”, allowing us to construct the first Beveridge curve for Canada. We find that the calibrated matching model does a good job in reproducing the time-series of unemployment and to a lesser extent non-participation. Our results also indicate that allowing for duration dependence in flows between unemployment and non-participation is crucial for explaining overall levels in long-term joblessness observed in the data, and we find that changes in the duration distribution among unemployed and non-participants contributed less to the deterioration of labor market conditions in Canada, relative to the U.S. In part, this difference comes from the fact that the U.S. recession was much more severe.

*Kory Kroft, University of Toronto, Department of Economics and NBER (kory.kroft@utoronto.edu); Fabian Lange, McGill University, Department of Economics, Leacock Building, 855 Sherbrooke Street West, Montreal Quebec H3A2T7, Canada (fabolange@gmail.com); Matthew J. Notowidigdo, Northwestern University, Department of Economics and NBER (noto@northwestern.edu); Matthew Tudball, University of Toronto (matthew.tudball@utoronto.edu). The authors would like to thank Robert French and Susan Ou for excellent research assistance and David Card, Philip Oreopoulos, and participants at the Ottawa, October 2016 “Small Differences” conference for helpful comments. We also thank the Bristol Workshop on Economic Policy Interventions and Economic Behavior as well as Samuel Bentolila. Much of the analysis for the paper was conducted at the Toronto Region Statistics Canada Research Data Centre, which is part of the Canadian Research Data Centre Network (CRDCN). The services and activities provided by the CRDCN are made possible by the financial or in-kind support of the SSHRC, the CIHR, the CFI, Statistics Canada and participation universities whose support is gratefully acknowledged. The views expressed in this paper do not necessarily represent the funders, CRDCN or any of its partners.

1 Introduction

The textbook model of the labor market features a matching function mapping unmet search demand and supply into new employment relationships (Mortensen and Pissarides [1994]; Pissarides [1985]). In previous work, we examined how this framework performed over the Great Recession in the United States in accounting for the dramatic rise in long-term unemployment and the outward shift in the Beveridge curve (Kroft et al. [2016]). Our main results showed that a standard matching model with unemployment (U) and vacancies (V) and a constant job-finding rate did a poor job fitting the data. We thus enriched this matching model along two dimensions. First, we allowed for duration dependence in the job-finding rate of the unemployed, consistent with a range of empirical evidence (e.g., experimental work such as Kroft et al. [2013] and recent structural econometric work such as Bentolila et al. [2017]). Second, we allowed for flows between employment (E), unemployment (U), and non-participation (N), instead of focusing exclusively on flows between employment and unemployment, as in the standard matching model.¹ We calibrated our enriched matching model and found that it could account for most the rise in long-term unemployment and about half of the shift in the Beveridge curve.

In this paper, we compare the labor market dynamics in the U.S. and Canada before, during, and after the Great Recession, building on the model in Kroft et al. [2016] and using restricted-use panel from the Canadian Labor Force Survey (LFS) – the counterpart to the U.S. Current Population Survey (CPS). This comparison is of interest because, as we demonstrate, the dynamics of the Great Recession in Canada and the U.S. were quite different. By focusing on Canada, we thus subject our matching model to a new “out-of-sample” test. In the process, we also expand and build on our prior analysis in several ways to make four contributions.

First, we document how the U.S. labor market has evolved since March 2013 (where our prior analysis left off). Our results indicate that the model continues to explain very well the share of long-term unemployment as well as the outward shift of the Beveridge curve. However, the model continues to over-estimate the job-finding rate among non-participants and thus under-estimates the stock of non-participants. This continues to remain one of the enduring puzzles of the Great Recession in the U.S.

Second, we present new evidence on unemployment dynamics in Canada during the Great Recession and compare it to the U.S. We document a rise in long-term unemployment in Canada that was less pronounced than the rise in the U.S. Similar to the U.S. (see Kroft et al. [2016]), we show that observables cannot explain the rise in long-term unemployment; rather, the increase was widespread and occurred within many different demographic groups. We also find that at the onset of the recession job-finding rates

¹Throughout the paper, we will use non-participants (N) and out-of-labor force (OLF) interchangeably.

among the unemployed declined by similar amounts in the U.S. and in Canada. However, the decline in the job-finding rate in the U.S. persisted for much longer. Our results also indicate that Canada did not experience a decline in the rate at which non-participants transitioned back into employment, contrary to our findings for the U.S.

Third, we exploit a unique feature in the LFS data which allows us to augment and improve the matching model in Kroft et al. [2016]. In particular, the LFS asks both the unemployed and non-participants about their (ongoing) length of unemployment as well as their (ongoing) length of joblessness.² We refer to the duration of time since the most recent job as “joblessness duration”. Using these data on joblessness durations allows us to extend our model in a straightforward yet important direction to allow for duration dependence in the job-finding rates of both the unemployed and non-participants and in the flows between unemployment and non-participation.

We find that allowing for duration dependence in all flows involving non-participants and unemployed is important for matching the level and dynamics of long-term joblessness in Canada. In particular, we find that the unemployed become more likely to transition to non-participation and non-participants become less likely to transition to unemployment as the duration of joblessness increases. We also observe that job-finding rates for non-participants decline with duration of joblessness, with the decline similar in magnitude to the one for the unemployed. It is impossible to capture the levels and dynamics of long-term joblessness in Canada without considering these patterns in duration dependence, particularly those in the flows between unemployment and non-participation. By contrast, duration dependence helps but is not central to understanding the aggregate rates of employment, unemployment, and non-participation observed in Canada over the 2008-2009 recession. The observed behavior of these stocks is dominated by average flows between the three labor force states, which can be adequately captured without modeling long-term joblessness.

Our fourth contribution is to construct a new vacancy series for Canada which we use to develop a Beveridge curve for Canada. This is necessary because (to our knowledge) there is no widely-accepted

²There are a priori reasons why one might want to use measures of joblessness duration as opposed to unemployment duration. As noted by Clark and Summers (1979), joblessness duration is a much more useful measure than unemployment duration in order to understand the costs of unemployment. Joblessness duration, as opposed to unemployment duration, is also conceptually easier for survey respondents to understand since it does not require actively tracking time spent searching for a job. Kudlyak and Lange [2017] document that the duration data based on the survey responses in the CPS is inconsistent with the observed patterns of unemployment, non-participation, and employment observed in the CPS. In many instances, individuals report long durations of unemployment even though they were just observed to transition from OLF or employment to unemployment. In part, these responses are in fact consistent with the questions asked by the CPS. The duration question in the CPS refers to time spent searching for employment. Respondents may report long durations even though they just transitioned from employment into non-employed because they continued searching while holding stop-gap jobs. That is, the CPS-respondents might often give logically consistent answers to the questions they are asked, while labor market researchers make a logical leap in interpreting these data as representing unemployment durations. Kudlyak and Lange [2017] also present evidence from the 4-month CPS panels in the U.S. that suggests the same patterns of how joblessness affects transitions that we observe in Canada are also present in the U.S. In particular, they show that job-finding rates conditional on non-participation decline rapidly in the duration since last employment.

vacancy series in Canada that spans the years before, during, and after the Great Recession. Our vacancy measure is based on relative employment in “recruiting” industries and is inspired by Landais et al. [2016]. Landais et al. [2016] use employment in recruiting industries to create a reliable proxy for the U.S. where vacancy data are more widely-available and so the proxy can be validated. We follow the same approach to create a proxy for the vacancy rate in Canada. We use this proxy to calibrate the Canadian matching model and to create a Beveridge curve for Canada, which may be of independent interest. We demonstrate that the Beveridge curves in Canada and the U.S. have a broadly similar shape. However, unlike the U.S., the Beveridge curve in Canada did not shift outward significantly during the Great Recession.

The remainder of the paper proceeds as follows. In Section 2 we describe the data sources used in the analysis and the construction of our vacancy proxy using “recruiting” industries. In Section 3, we compare and contrast the recessions in Canada and the U.S. Section 4 revisits the analysis in Kroft et al. [2016] using U.S. data, replicating the main results and extending the results to the most recent data available. Section 5 describes patterns of incidence of long-term joblessness among the unemployed in Canada and reports results from composition analysis. Section 6 describes the extended model and discusses the model calibration and counterfactual results using the restricted-use LFS data. Section 7 concludes.

2 Data

This section briefly describes our data sources. See Kroft et al. [2016] for more details on the U.S. data.

2.1 United States

Current Population Survey (CPS)

We use monthly CPS data between 2002 and 2015 (ending in October 2015, extending beyond the April 2013 end date in Kroft et al. [2016]). Our sample comprises all employed, unemployed, and non-participants respondents aged 25-55. The data is age-adjusted to the 2000 age-distribution in the U.S. In the cross-section, we keep track of the total population of each category to estimate the “stocks”. To create panel data, we match observations across successive months, matching on household identifier, line number, age, gender, and race. We use the matched panel data in addition to the CPS cross-sectional estimates of the unemployed, the employed, and non-participants to estimate the transition rates between unemployment, employment, and non-participation in each month. We also compute overall (pre-2008) transition rates by unemployment duration (into both employment and non-participation). Finally, we

compute transition rates from employment and non-participation into unemployment by unemployment duration.

Job Openings and Labor Turnover Survey (JOLTS)

We use monthly JOLTS data between 2002 and 2015 to compute the total number of vacancies. We use these vacancy data to calibrate the matching model below during the pre-2008 period. We then use the post-2008 vacancy data as one of the exogenous "forcing variables" for our counterfactual scenarios.

2.2 Canada

Labor Force Survey (LFS)

We use restricted-access Canadian Labor Force Survey (LFS) data. The definition of unemployment varies somewhat between the U.S. and Canada. To be able to compare the data from the two countries, we went back to the survey responses to the LFS and recoded labor force states in the LFS to be comparable to those from the CPS (see Bernard and Usalca [2014] and references contain therein). Similar to the U.S., we limit ourselves to the age range 25-55 and re-weight the data to match the 2000 age distribution. We also force the transition rates to be consistent with the cross-sectional data on stocks using the methodology described in Kroft et al. [2016].

An important advantage of the Canadian data is that respondents are asked about the duration of joblessness in addition to the duration of unemployment. As Elsby et al. [2013a] and Kroft et al. [2016] both report, unemployment durations as usually constructed based on the CPS are frequently inconsistent with the observed panel of labor force states when observations in the CPS are linked to construct an individual-level panel data set. In the CPS, unemployed respondents are asked to report how long they have been actively looking for a job. It has become standard practice to use reported durations in response to that question as indicating unemployment durations. If respondents interpreted this question as labor statisticians would like them to, they should report durations of less than a month when surveyed in the month after a transition from either employment or out of the labor force to unemployment. However, respondents often report substantially longer durations of unemployment during these months. Kudlyak and Lange [2017] present evidence that respondents might include short jobs (stop-gap jobs) during which they continued to search for more permanent employment when they answered this question. By contrast, the question on the duration of joblessness in the LFS is relatively simple to interpret. This suggests that responses to the duration of joblessness are more likely to align with the concept in question. In addition, the duration of joblessness question is asked both of the unemployed and of non-participants.

We therefore focus on joblessness duration for the bulk of our empirical analysis, both for conceptual reasons and to have same duration concept for both unemployed individuals and non-participants.

2.3 A New Vacancy Series for Canada

In Kroft et al. [2016], we used the Job Openings and Labor Turnover Survey (JOLTS) to compute the total number of vacancies each month. Unfortunately, to our knowledge there is no counterpart data series in Canada that allows us to directly compute vacancies for the relevant time period. The vacancy measure from the Job Vacancy Statistics (JVS) series is only available since 2011 and the measure from the Job Vacancy and Wage Survey (JVWS) is only available since 2015.³ While other data sources exist, none is adequate for our purposes.⁴

Recently, Landais et al. [2016] developed a proxy for vacancies in the U.S. This proxy is called the "recruiting-producer ratio" and is defined theoretically as the ratio of the number of recruiters relative to the numbers of workers engaged in production. Empirically, Landais et al. [2016] define this ratio as:

$$\tau = \frac{\xi \times rec}{l - \xi \times rec}$$

where *rec* is the seasonally-adjusted monthly number of workers in "recruiting industries", defined as employment in North American Industry Classification System (NAICS) code 56131, and *l* is the seasonally-adjusted monthly number of workers in all private industries.⁵ The parameter ξ is a scaling factor used to adjust for labor devoted to recruiting by firms not belonging to the recruiting industry and therefore not captured by *rec*. In the U.S., Landais et al. [2016] set $\xi = 8.4$ based on survey evidence from 1997. Why should this measure be correlated with vacancies? The basic idea is that when firms are posting relatively more jobs, there are more resources tied up in recruiting. Thus, τ should be procyclical, and this is indeed what Landais et al. [2016] find.

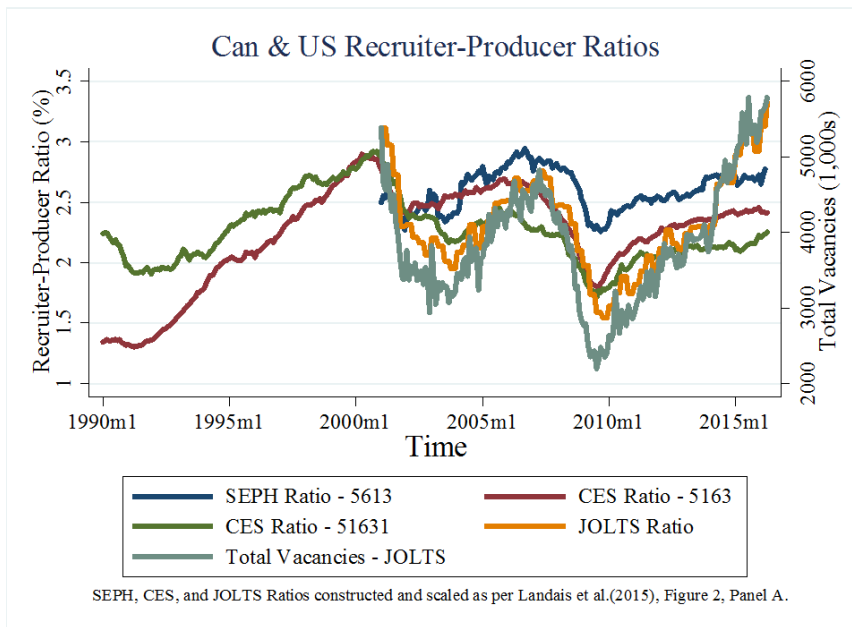
We follow Landais et al. [2016] and construct a vacancy series using employment in recruiting industries. Unfortunately, employment counts by 5-digit NAICS codes are not available in Canada. Thus, we instead use employment at the 4-digit NAICS industry code 5613 measured in the Survey of Employment,

³The JVS is reported monthly and the data is collected at the establishment level, whereas the JVWS is reported quarterly and the data is collected at the business location level. The JVWS also has a larger sample, and includes vacancies from businesses primarily involved in agriculture, which the JVS does not. For more information on these two vacancy series, see the publication <http://www.statcan.gc.ca/pub/75-514-g/75-514-g2015002-eng.htm>.

⁴The Canadian Federation of Independent Businesses (CFIB) produces a survey of vacancies. However, it suffers from a number of limitations. First, it excludes the public and utilities sector. Second, it allows for passive job search and doesn't require an open position to exist, just for the business to have an unmet need. Third, it comes from registered CFIB members who voluntarily take the "Your Business Outlook Survey". Non-responses to the vacancy question are coded as zero vacancies. The other source of vacancy data comes from the Conference Board which uses online data from roughly 80 job-posting websites collected by WANTED Technologies. However, this vacancy series also suffer from serious limitations and was in fact discontinued because of reliability issues.

⁵This official name for this industry is "employment placement agencies and executive search services".

Figure 1: Canadian & U.S. Recruiter Producer Ratios



Payrolls and Hours (SEPH) to estimate the number of employees employed in the recruiting industries.

To construct τ , we also need an estimate of ξ , the proportionality factor meant to capture the ratio of workers outside the recruiting industry that are engaged in recruiting. We expect our estimate ξ to be smaller than the estimate from Landais et al. [2016] since we define the recruiting industry to be broader.⁶ Below, we will estimate the adjustment factor ξ as one of the parameters entering the matching function. The estimates will depend on the model specification and in particular on whether we allow for duration dependence in the job-finding rates of the unemployed and/or those OLF. Our preferred specification allows for both duration dependence among the unemployed and those OLF. The estimate of ξ for our preferred specification is 3.5 and we will use this value when we describe the vacancy series. Fortunately, the measures τ and its cyclical properties are robust to variation in ξ and even substantial variation in this parameter generates very little differences in the time-series of τ .

Figure 1 compares the recruiter-producer ratio in the U.S. with the vacancy measure taken from JOLTS and also shows the recruiter-producer ratio for Canada. Note that in the U.S., the recruiter producer ratios based on NAICS code 56131 and 5613 behave very similarly over time, suggesting that basing our vacancy series on the 4-digit NAICS code in Canada will not be a major problem for our analysis. Furthermore, over time, the recruiter-producer ratio and JOLTS show very similar time series

⁶The seasonally-adjusted employment counts l stem from CANSIM table 281-0047.

behavior in the U.S. - rapid declines at the onset of the 2001 and 2008 recessions, with recovery in between and comparable relative movements during the recovery and the recessions. To provide another way of validating the vacancy proxy in the U.S., in the Appendix we report predicted job-finding rates (for both unemployed and non-participants) in the U.S. using both the JOLTS measure as well as the recruiter-producer ratio proxy. The predicted job-finding rates are very similar across the two measures, bolstering the case that the proxy based on the recruiter-producer ratio in Canada may be an adequate substitute for JOLTS-like vacancy series which currently do not exist for a sufficiently long time-period in Canada.

Comparing the U.S. and the Canadian time-series patterns in vacancies as measured by the recruiter-producer ratio, we find that over the last recession, the recruiter-producer ratio declined much less in Canada than in the U.S. This finding lines up with the general observation that the recession was less severe in Canada than the Great Recession in the U.S.; Canada experienced a “Not-Quite-Great Recession”.

3 The Recessions in Canada and the U.S.: A Brief Overview

3.1 Unemployment and E-Pop Rates

Both the U.S. and Canada experienced a rapid, sharp increase in the unemployment rate during the Great Recession, but the magnitude, persistence, and onset of the Great Recession differed between the two countries. The NBER has determined that the U.S. was in a recession for 18 months from December 2007 to June 2009, while the recession in Canada lasted only 7 months from November 2008 to May 2009.⁷ The recessions differed not just in length but also in their severity. The movements in key labor market outcomes, illustrated in figures 2 and 3, were about twice as large in the U.S. as in Canada. Figure 2 shows that for the U.S., the unemployment rate among 25-55 year olds increased by about 5 percentage points over the Great recession compared to a 2.5 percentage point increase in the unemployment rate in Canada. In the U.S., the employment-to-population ratio (figure 3) declined by about 4 percentage points compared to a 2 percentage point decline in Canada.⁸

The figures also show that Canada recovered more rapidly than the U.S. By mid-2010, the unemployment rate in Canada had dropped back by about half of its increase in 2009, whereas for the U.S., the unemployment rate remained at levels similar to those in late 2009. Since then, both the U.S. and

⁷See <http://www.nber.org/cycles/cyclesmain.html> for the U.S. and <https://www.cdhowe.org/cpc-communique/cd-howe-institute-business-cycle-council-issues-authoritative-dates-20082009-recession> for Canada.

⁸In this paper, all statistics are constructed using prime age populations (25-55) and are age adjusted to the 2000 age distributions in the two countries. The duration concept in Figure 4 is the duration of unemployment (not joblessness).

Figure 2: Unemployment in the U.S. and Canada 1/2001-10/2015

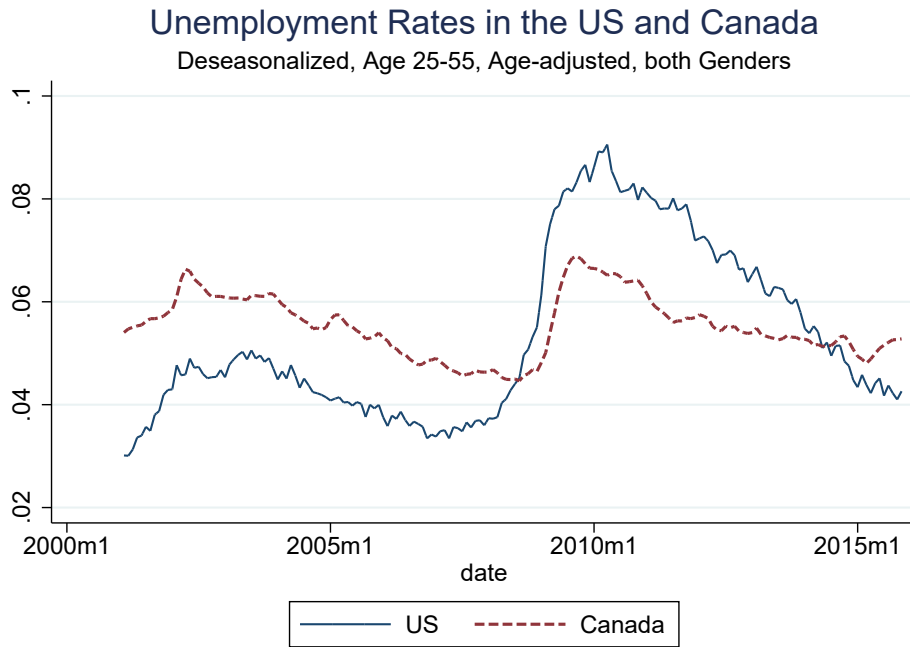


Figure 3: Employment to Population Ratio: U.S. and Canada 1/2001 - 10/2015

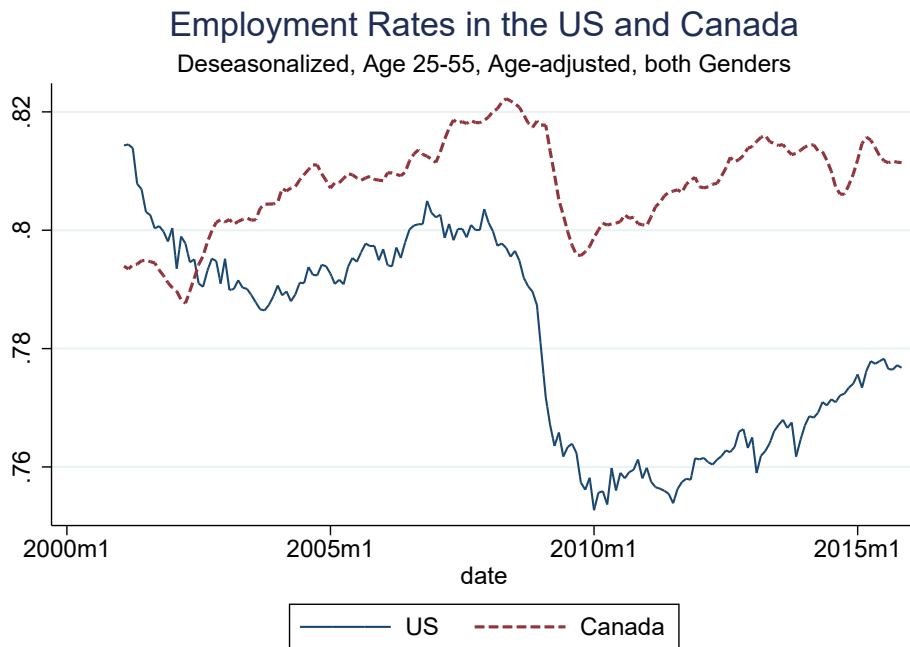
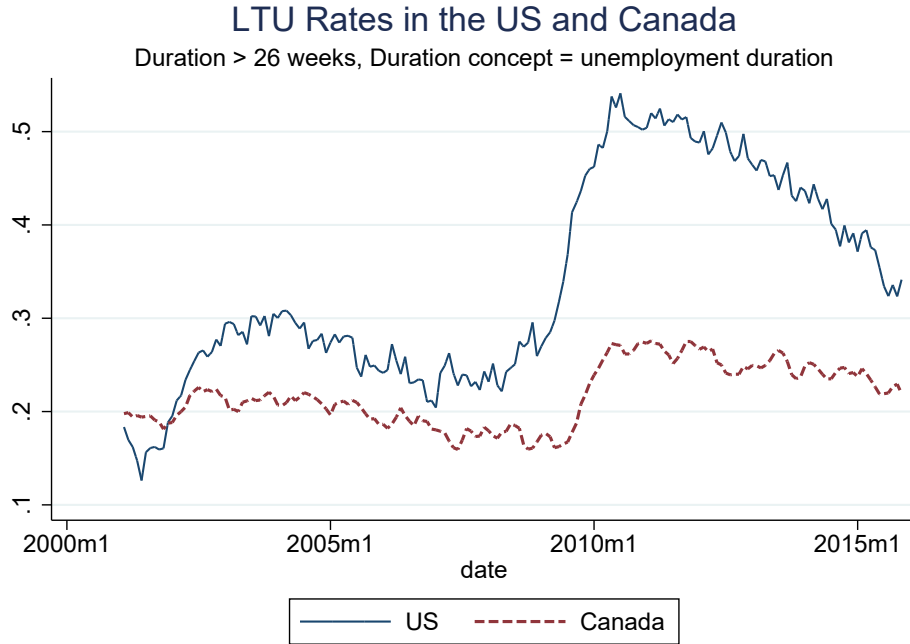


Figure 4: LTU U.S. and Canada 1/2001 - 10/2015



Canada slowly returned to unemployment rates comparable to the pre-recession period. However, the U.S. has only managed to claw back about half of the decrease in the employment-to-population ratio, while Canada returned to the pre-recession levels of the employment-to-population ratio by the end of 2012.

3.2 Long-Term Unemployment and Joblessness

In addition to the overall increase in unemployment, the Great Recession also increased unemployment at long durations. Figure 4 shows the share of the unemployed with unemployment durations exceeding 6 months for the U.S. and the share of the unemployed with durations exceeding 26 weeks for Canada. In both countries, the long-term unemployment (LTU) share increased significantly, but the rise was much more pronounced in the U.S. In the U.S., the LTU share increased by about 25 percentage points, while in Canada it rose by only about 10 percentage points. Additionally, the decline in the LTU shares in both countries has been very slow. By October 2015, the LTU share in the U.S. was about 30 percent and in Canada was about 20 percent.

As described above, a useful feature of the Canadian data is that we can observe *joblessness* durations for both the unemployed and non-participants, which is something that the CPS does not keep track of. Figure 5 shows the share of long-term joblessness (LTJ) for the unemployed. We see that it fluctuates

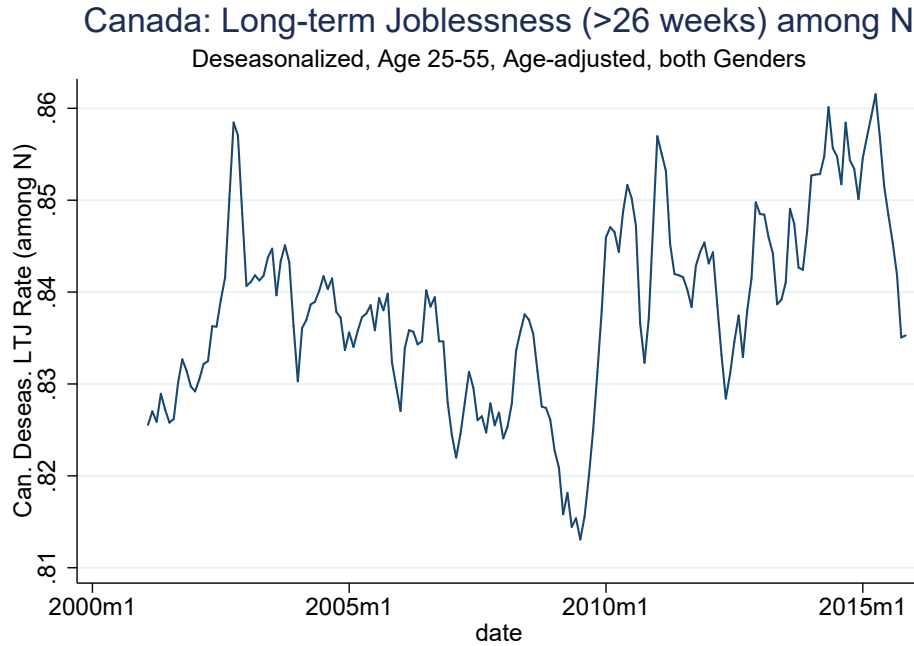
Figure 5: Long-Term Joblessness among the Unemployed in Canada



in the pre-recession period around 45 percent and increases to roughly 55 percent in the post-recession period. Among the unemployed, the increase in percentage points in long-term joblessness is about of the same order of magnitude as the increase in long-term unemployment. In percent, it is much smaller since the base incidence of long-term joblessness is substantially larger.

Figure 6 shows the share of long-term joblessness for non-participants. Rates of long-term joblessness are much higher for non-participants than they are for the unemployed. This reflects the fact that this group is composed of many individuals who are not on the margin of entering the labor market. Similar to the trends for the unemployed, the LTJ share among non-participants declined prior to the recession and increased moderately during the recession. While most of the increase in long-term joblessness concentrated in the first few months of the Great Recession among the unemployed, the increase in LTJ among non-participants has been much more gradual. It is indeed not clear whether this increase in LTJ has run its course by October 2015, when our data ends. Overall, the relative increase in LTJ is much more pronounced among the unemployed since the share of LTJ is of course much lower among the unemployed. Nevertheless, since there are many more non-participants as compared to unemployed, an increase of 4 percentage points in the LTJ share among non-participants is an important empirical pattern.

Figure 6: LTJ among OLF



3.3 Transition Rates

Figure 7 provides another way of examining the relative labor market performance in the U.S. and Canada over time. It shows the job-finding rates and the job-loss rates depending on whether they involve unemployment or non-participation. The data is very noisy so we present 6-month moving averages centering on the current month. The top two panels show the job-finding rates conditional on U or N and the bottom two panels show the rates at which the employed transition into U or N.

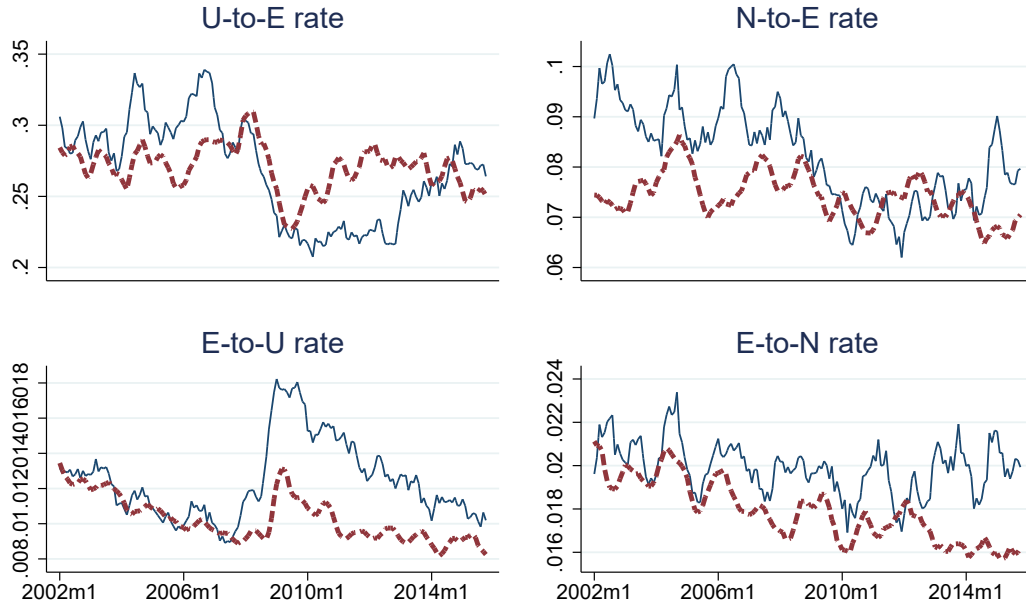
Starting with the job loss rates, we see that for Canada, the E-to-U rate briefly increased from 1 percent to about 1.5 percent, but quickly returned back to normal levels. By contrast, in the U.S., the E-to-U rate increased from about 1 percent to around 2 percent and remained elevated during the Great Recession. In both countries, the rates at which individuals transition from employment to non-participation were relatively stable. Considering the job-finding rates, we observe an initial decline for the unemployed that is similar in size in both Canada and the U.S. However, we again observe that Canada returned to its prior levels more rapidly. Among non-participants, unlike the U.S. which experienced a drop in the job-finding rate, we do not observe a significant decline in job-finding rates in Canada. It is striking that in the U.S., job-finding rates among non-participants so far have not fully recovered to their pre-recession levels.

It is not the initial decline in job-finding rates among the unemployed that explains why the U.S.

Figure 7: Job Loss and job-finding Rates in the U.S. and Canada

US and Canada: Job Finding and Job Loss Rates

US = solid, thin line, Canada = thick, dashed line



Age 25-55, age-adjusted to 2001 distribution, Deseasonalized. 6 month moving average.

recession was “Great” and that in Canada was not. Rather it is the persistence in the declines in job-finding rates and job loss rates in the U.S. that transformed the U.S. recession into a Great recession. It seems that Canada escaped a Great Recession primarily because its labor market rebounded more quickly and because job-finding rates among those OLF did not deteriorate substantially.

3.4 Beveridge Curve

The Beveridge curve is one of the main diagnostic tools used to understand the labor market performance over the business cycle. Up to recently, vacancy series were not available for Canada, making it impossible to construct a Beveridge curve describing the Canadian labor market performance. Using the vacancy series constructed according to the recruiter-producer ratio we are in a position to produce what we believe to be the first Beveridge curve for Canada.

Figure 8 shows this Beveridge curve depicting the vacancy to population rate and the unemployment to population rate in our data. The data depicted here are quarterly averages running from Q1 2001 to Q3 2015 and clearly show the downward-sloping relationship familiar from Beveridge curves drawn in other countries such as the U.S. or the UK. Compared to the U.S., we do not see a shift in the Beveridge curve

Figure 8: The Beveridge Curve in Canada

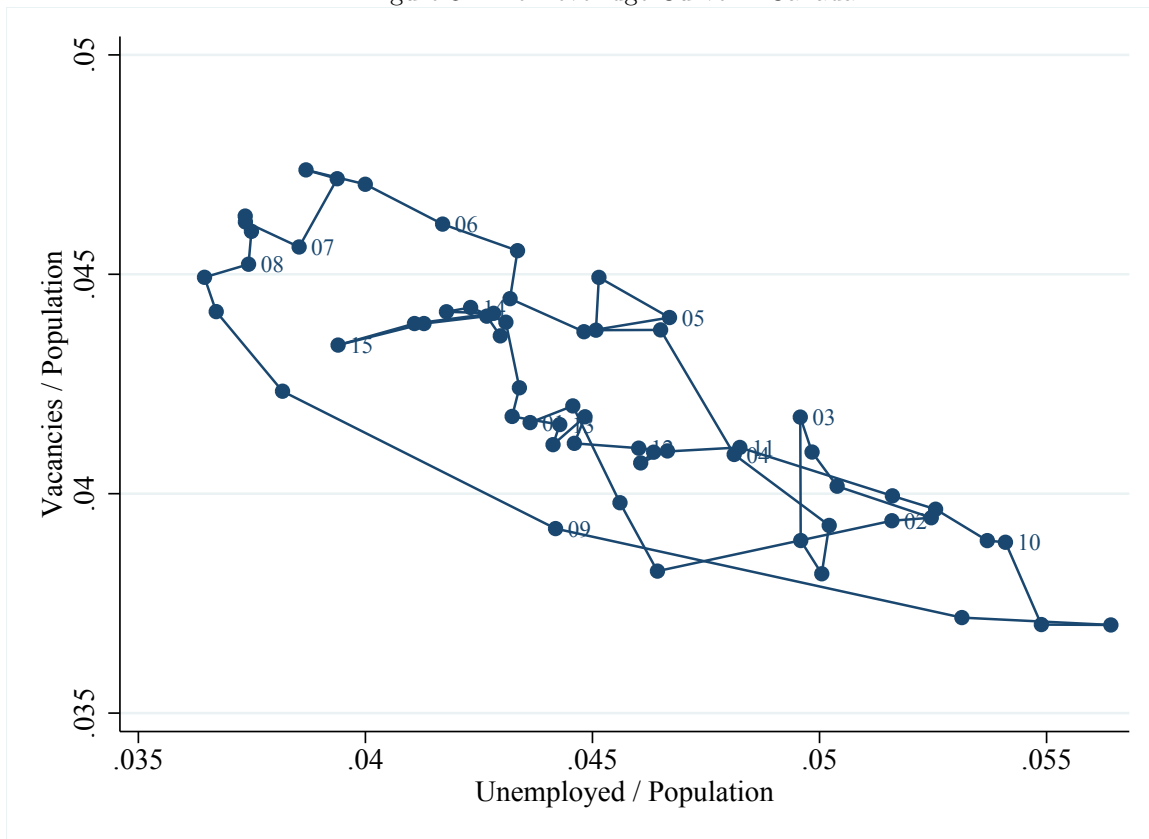
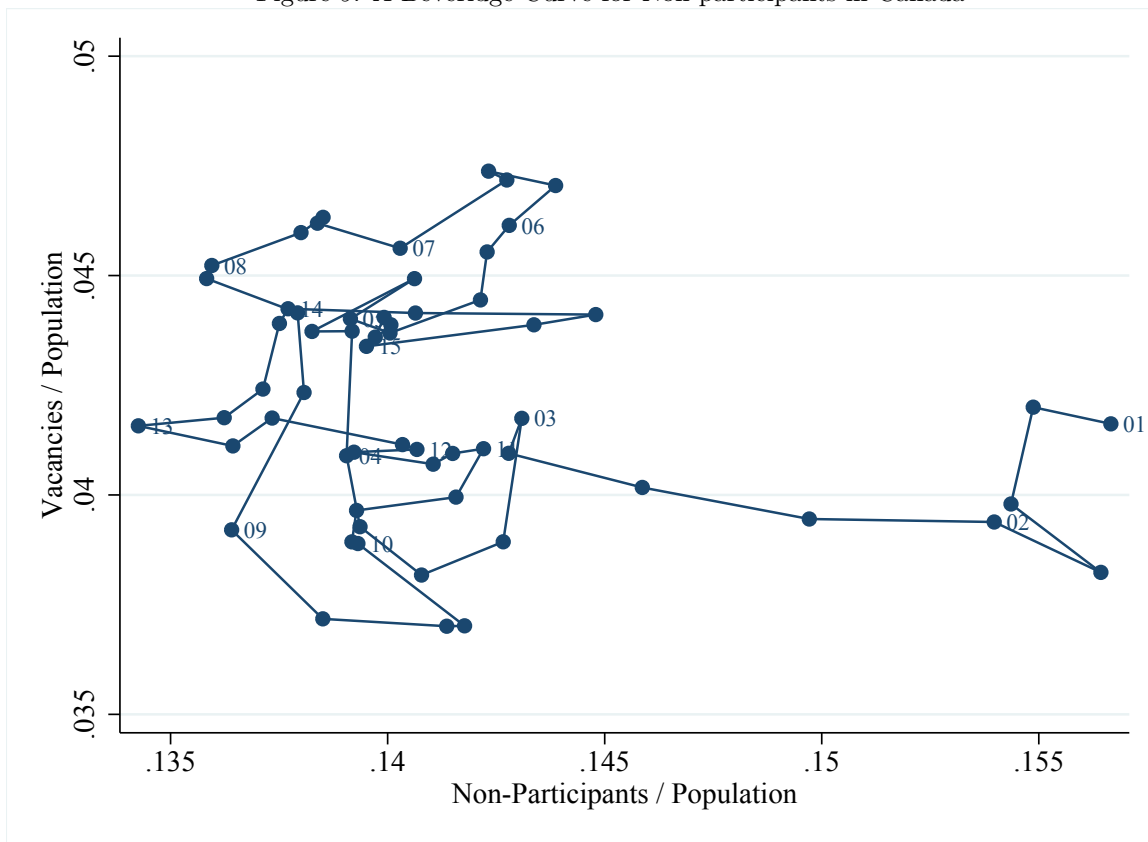


Figure 9: A Beveridge Curve for Non-participants in Canada

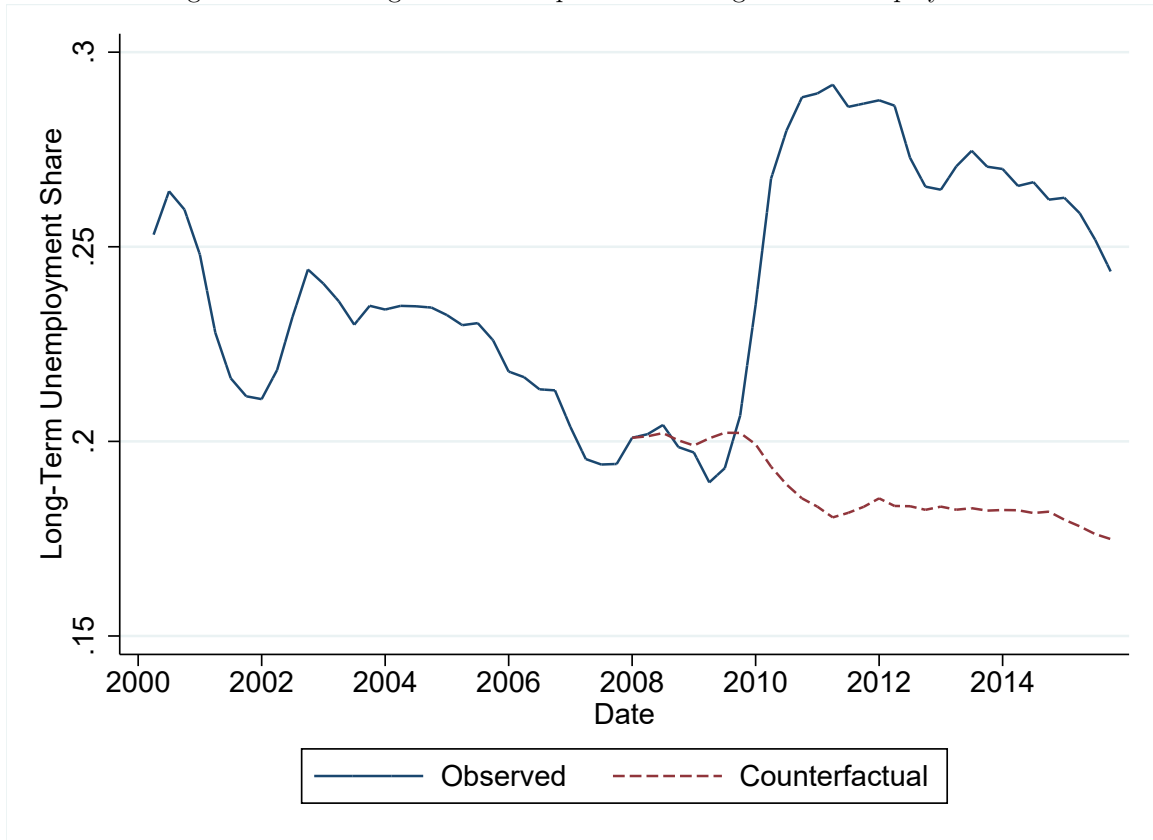


over the Great Recession in Canada. We also show the analogous relationship between vacancies and non-participants in Figure 9. It is clear that there is less of a systematic relationship between vacancies and non-participants than between vacancies and unemployment.

4 The Role of Composition for Trends in Long-Term Unemployment and Long-Term Joblessness in Canada

As described above (see Figures 5 and 6), the rate of long-term unemployment and long-term joblessness in Canada rose sharply at the end of 2009 and has since remained elevated. Roughly half of the increase still remains. We now explore this increase in LTU and LTJ in more depth. We first assess how much of the growth in LTU and LTJ in Canada can be accounted for by changes in the observable composition of the unemployed and non-participants. Specifically, we consider education (high school dropout, high school grad, some college, and college graduate), age (6 5-year age groups between 25 and 55), and region, and gender. We proceed to separately investigate the role of the composition across these categories in the time-patterns of the LTU among the unemployed and LTJ among non-participants. Long-term

Figure 10: Assessing Role of Composition in Long-Term Unemployment



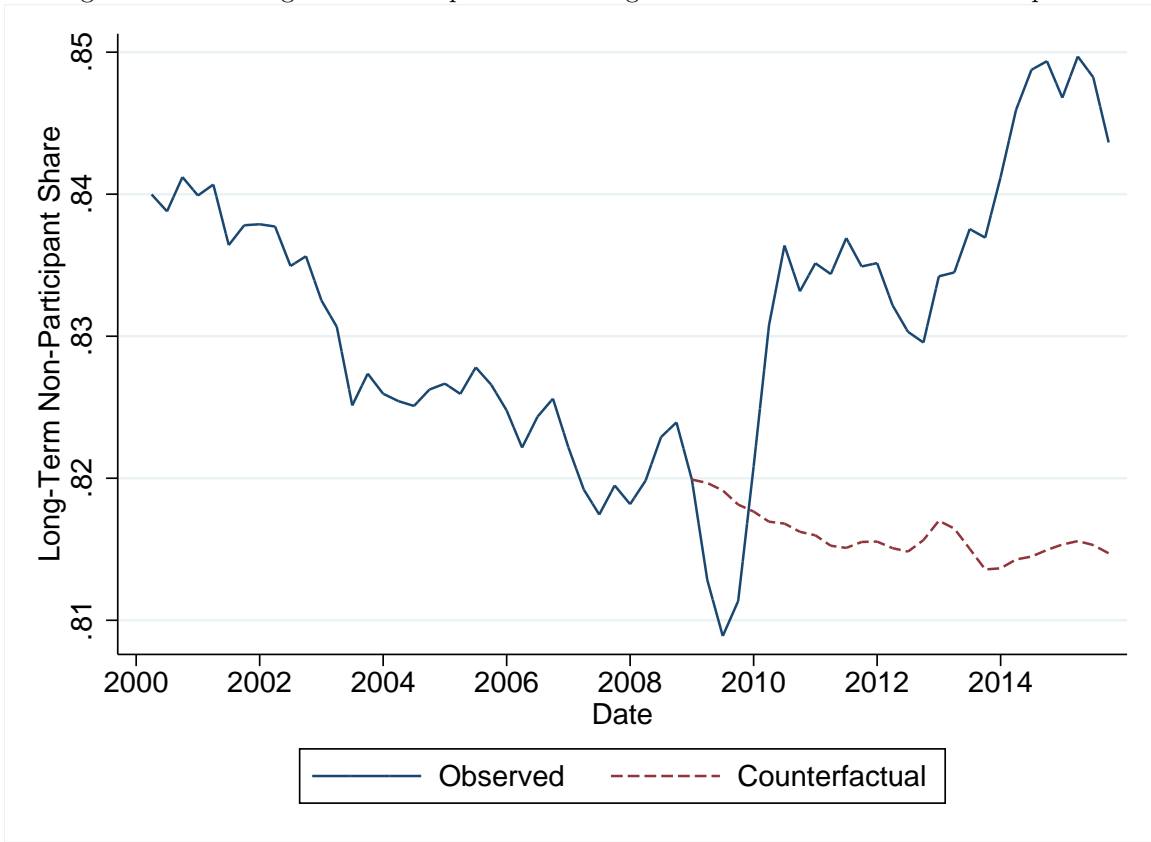
unemployment and long-term joblessness are both defined relative to a 6 month cut-off.

In the Appendix, we present both the incidence of LTU and the share conditional on each of the above listed characteristics among the unemployed. These figures illustrate that the variation in LTU over time is quite similar across different characteristics. Furthermore, when there is variation in LTU across observable characteristics, the share of the different groups among the unemployed does not change dramatically when the recession hits. Consequently, there is limited scope for changes in the composition of unemployed to account for overall change in LTU for unemployed workers.

Figure 10 combines all of these characteristics together to predict the change in LTU from compositional changes in the sample of unemployed individuals using a shift-share procedure. We find that the predicted change in the overall LTU share based on compositional changes is very small, mirroring the results we found in Kroft et al. [2016] for the U.S. To create this figure, we create linear prediction of LTU based on observable characteristics, using pre-2007 incidence of LTU across different characteristics. Our projection uses each of the characteristics shown in Appendix.

We can also repeat the same exercise for non-participants. In the Appendix, we show a similar pattern of long-term joblessness across education groups, but large differences in levels. In addition to the large

Figure 11: Assessing Role of Composition in Long-Term Joblessness for Non-Participants



pre-recession differences in LTJ levels by education, there are also meaningfully different trends in the composition of long-term jobless non-participants, with more long-term jobless non-participants having high education in recent years (as compared to earlier years). Despite these differences in levels and trends by education, however, the limited overall role of compositional changes in accounting for LTJ trends among non-participants is shown in Figure 11. As with the long-term unemployed in both U.S. and Canada, the results suggest no meaningful role for compositional changes in accounting for national trends in long-term joblessness.

Thus, overall the rise in LTU and LTJ in Canada after 2009 was not isolated to specific demographic groups, but rather was experienced broadly across the labor market. These results using observables do of course not speak directly to the potential for compositional changes based on unobservables. Nevertheless, they do suggest that changes in the composition of the unemployed are not driving the observed patterns in LTJ during the recession.

5 Replicating and Extending Kroft et al. [2016]

So far, we have shown how the labor market experience in Canada during the last recession differed from the U.S. experience. For the remainder of the paper, we explore the ability of the standard matching model augmented with duration dependence in labor market flows to match the observed patterns in unemployment and non-participation as well as long-term unemployment and joblessness over the business cycle. This work builds substantially on Kroft et al. [2016] and we begin by reviewing the methodology employed in that paper.⁹ Our treatment here is sparse and we refer the reader to Kroft et al. [2016] for details.

The Matching Function

At the core of our analysis of the U.S. labor market in Kroft et al. [2016] is a matching function which determines the number of meetings between job openings and both unemployed and non-participants.

$$M(U_t + sN_t, V_t) = m_0 (U_t + sN_t)^\alpha V_t^{1-\alpha} \quad (1)$$

Here (U_t, N_t) are the unemployed and non-participants; V_t are vacancies and (m_0, s, α) are parameters. One can interpret $(U_t + sN_t)$ as the total units of search effort on the labor supply side where each unit of search effort delivers an identical probability of a meeting with a vacancy. Then s represents the relative search effort of those deemed OLF. The probability of a meeting per unit of search effort is $\frac{M(U_t + sN_t, V_t)}{U_t + sN_t} = m_0 x_t^{1-\alpha}$ where $x_t = \frac{V_t}{U_t + sN_t}$ is a measure of market tightness that accounts for the non-participants.¹⁰

The function $A(d)$ is defined as the relative job-finding rate of unemployed of different durations of unemployment d . This function captures “true” duration dependence (sometimes called “structural duration dependence”); i.e., the genuine causal effect of longer durations on the job-finding rate. As described in Kroft et al. [2016], this modeling assumption is motivated partly by recent field experimental evidence on duration dependence in “callbacks” for interviews (Ghayad [2013], Kroft et al. [2013]). It is also consistent with recent structural econometric work that finds evidence of “true” duration dependence using population-level data from Spain (Bentolila et al. [2017]). This assumption is also consistent with recent work finding a causal effect of duration of joblessness on re-employment wages (Nekoei and Weber [2017], Schmieder et al. [2016]), since these papers interpret these results as suggesting human capital depreciation, which would cause “true” duration dependence. We emphasize that understanding how

⁹Some of the formulations in this section are lifted directly from that paper and edited for brevity.

¹⁰See Hornstein et al. [2014b] for an index of labor search that accumulates various different groups among the non-employed including the marginally attached, the discouraged and those with diverse labor search histories in a similar manner.

much of the $A(d)$ function that we estimated from observation data represents “true” duration dependence remains very much an open question, with a range of recent evidence suggesting a smaller role for “true” duration dependence in generating observed pattern of duration dependence (Farber et al. [2015], Ahn and Hamilton [2016], Jarosch and Pilossoph [2016]). In this paper, we will proceed under the assumption that we have access to the function $A(d)$ representing “true” duration dependence. We do however recognize that the extend of true duration dependence is very difficult to assess from observational data and thus explore sensitivity to alternative assumptions about $A(d)$ function. With $A(d)$ in hand, the job-finding rates of the unemployed with duration d and the non-participants in the model used to analyze the U.S. labor market in Kroft et al. [2016] are then given by

$$\lambda^{UE}(x_t; d) = A(d)m_0x_t^{1-\alpha} \quad (2)$$

$$\lambda^{NE}(x_t) = sm_0x_t^{1-\alpha} \quad (3)$$

We normalize $A(0) = 1$ and assume that $A(d)$ follows a double-exponential decay function. We estimate the parameters of $A(d)$ and the parameter s using data on job-finding rates by duration and non-participation from the period 2002-2007 preceding the Great Recession. Kroft et al. [2016] reports the parameter estimates governing the transitions into employment which are estimated using the data from the pre-period 2002-2007. The functional form and parameters of the functions $M(U + sN, V)$ and $A(d)$ and s are constant over time. Thus, when we examine the performance of our model over the Great Recession, we ask whether the search environment accounting for duration dependence is stable over the Great Recession except for the demand for labor.

Simulating labor flows

To simulate the labor market, we need to measure or model flows between all three labor force states (E, U, N). In addition, we need to model the duration distribution of the flows into unemployment. Above, we showed how to model the endogenous flows into employment: $\lambda_t^{UE}(d)$ and λ_t^{NE} . The remaining flows in the labor market are exogenous processes in our analysis. As stated in Kroft et al. [2016], to construct these flows we assume that if “non-participants move to unemployment, they draw an unemployment duration from the (empirical) distribution of unemployment durations estimated from observed N-to-U transitions.” Similarly, when employed workers move into unemployment, they draw an unemployment duration from the empirical distribution of unemployment durations.¹¹ These two empirical distributions

¹¹One of the lessons we draw from the analysis in this paper as well as in Kroft et al. [2016] is that is crucial to account for the fact that transitions between unemployment and non-participation vary across the duration distribution. The assumptions made in Kroft et al. [2016] are ad-hoc, but they do succeed in roughly accounting for the fact that the

are defined as $\theta_t^{NU}(d)$ and $\theta_t^{EU}(d)$, respectively. The dynamic equations governing changes in the stocks are then

$$N_{t+1} = N_t \left(1 - \lambda_t^{NU} - \hat{\lambda}_t^{NE} \right) + E_t \lambda_t^{EN} + U_t \lambda_t^{UN} \quad (4)$$

$$U_{t+1}(0) = E_t \theta_t^{EU}(0) \lambda_t^{EU} + N_t \theta_t^{NU}(0) \lambda_t^{NU} \quad (5)$$

$$U_{t+1}(d) = U_t(d) \left(1 - \hat{\lambda}_t^{UE}(d) - \hat{\lambda}_t^{UN}(d) \right) + E_t \theta_t^{EU}(d) \lambda_t^{EU} + N_t \theta_t^{NU}(d) \lambda_t^{NU} \quad (6)$$

$$E_t = P_t - N_t - U_t \quad (7)$$

where P_t denotes the total population aged 25-55. We placed “ $\hat{}$ ” above the endogenous flow variables where as the other flow variables are exogenous driving variables (together with the vacancies) in our analysis.

5.1 Updating Kroft et al. [2016] to Oct. 2015

The analysis in Kroft et al. [2016] use data up until April 2013. In this Section, we explore the implications and performance of the model up since then.

Discrepancy from Kroft et al. [2016]

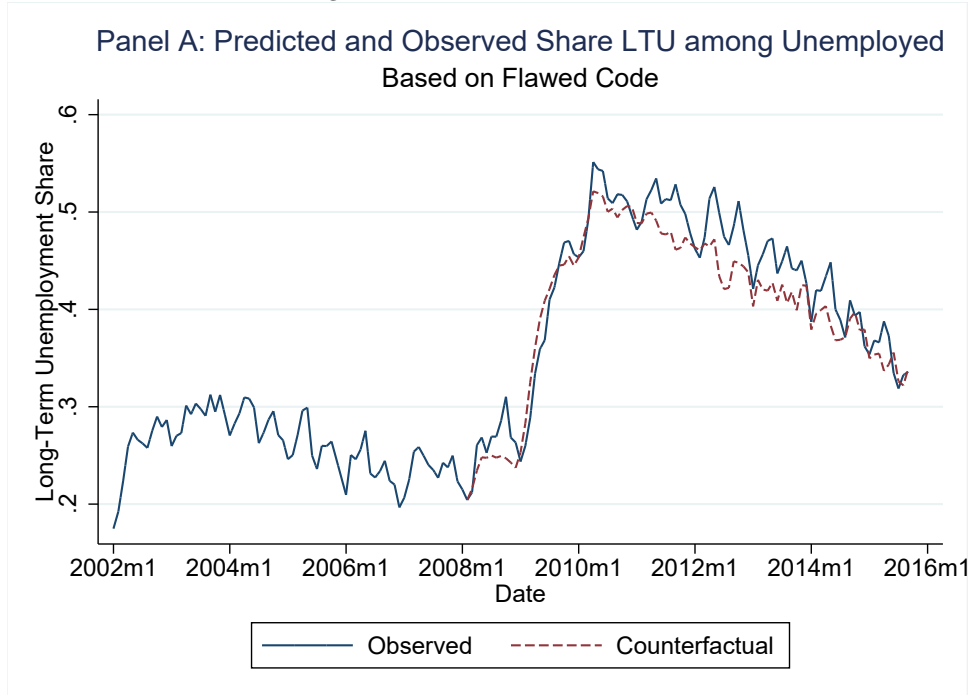
Before we turn to the updated analysis, we regret to report that some of the results shown here for the period up to April 2013 differ slightly from those reported in Kroft et al. [2016]. In preparation for this paper, we revisited the code used for Kroft et al. [2016] and discovered a coding error in the dynamic equations used in constructing the flows. In particular, the code did not exactly implement the equations above (which were written correctly in the paper). Upon re-running the analysis, we found that this mistake did not affect our basic conclusions but it did have a some effect on a few of the quantitative results.

In Figure 12, we compare the predicted and observed shares in LTU in Panel A and B. Panel A shows the results updated to October 2015 using our original code. By contrast, Panel B shows results using the updated code. This figure shows that we underestimate the actual share of long-term unemployed in the population of unemployed by about 5 to 10 percentage points.

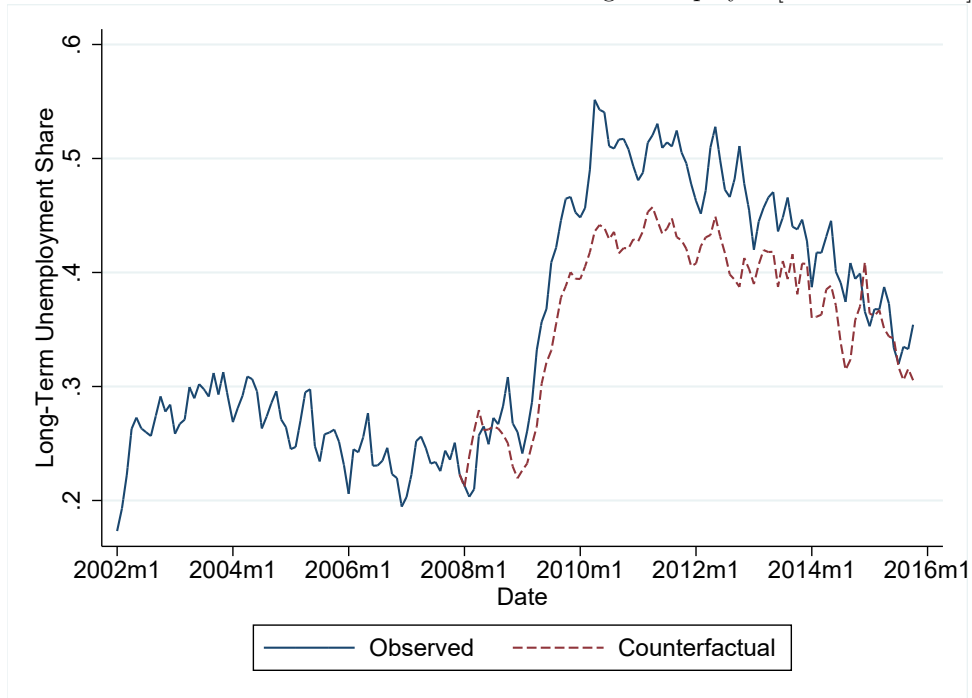
This implies that we reported a higher average duration of unemployment in our counterfactual series than warranted. By implication, our reported counterfactual job-finding rate for the unemployed was

durations reported by non-participants that transition to unemployment are relatively long. The advantage of the data on long-term joblessness for both U and N available in Canada is that we can replace these ad-hoc assumptions with direct measurements of the durations of those transitioning.

Figure 12: LTU Old and Fixed



Panel B: Predicted and Observed Share LTU Among Unemployed [Corrected Code]



too low. Consequently, while our prior simulations from Kroft et al. [2016] fit the job-finding rates conditional on unemployment very closely, we now tend to find a counterfactual job-finding rate conditional on unemployment that exceeds the actual job-finding rate by about 2 percentage points in the period following the Great Recession (see Panel A in 13). This tends to narrow as the discrepancy in the LTU share declines towards the later years in the sample.

The consequence of this coding discrepancy is fairly minor relative to the observed job-finding rates that typically vary between 20 and 30 percent. Given these large rates of finding jobs conditional on unemployment, a deviation of around 2 percentage points in the rate of finding jobs will not substantively affect the ability of the model to fit the simulated stocks in the labor market and thus leaves our conclusions largely unchanged. We still find that the model can account for a large share of the increase in LTU and some of the outward shift in the Beveridge curve, and we continue to find (as before) that the model does not do a good job accounting for trends in non-participation during and after the Great Recession.

Kroft et al. [2016] since April 2013

How did the model fare in the two-and-a-half years since we completed the previous analysis? We find that the model does quite well when considering the share of LTU (Figure 12) and the job-finding rates conditional on unemployment (Figure 13). Indeed, on both dimensions the model does better after April 2013. Figure 12 shows that the stock of LTU declines rather slowly but steadily. By the end of 2015, the observed and counterfactual share of LTU has declined to 35 and 30 percent respectively from about 50 and 40 percent respectively right after the Great Recession.

The model, however, fits the data much less well when we consider job-finding rates among the non-participants. Those rates are consistently overestimated after 2008 (see Panel B of Figure 13,). The persistent feature of the U.S. labor market that is most difficult to explain is not unemployment or long-term unemployment, but non-participation since 2008.

Consider next the Beveridge curve shown in Figure 14. Since April 2013, labor demand as measured by vacancies has continued to increase and unemployment continued to decline. However, even though vacancies exceed the rates observed at any time between 2002 and 2007, unemployment however has not declined to those levels seen in that period - indicating that the Beveridge curve has indeed shifted.

In Kroft et al. [2016] we reported that we could explain about half of the shift in the Beveridge curve until April 2013 using the changing duration structure. A gap of about 1 percentage point however remained to be explained outside of our model. Since then, we have seen the gap between the observed and the counterfactual Beveridge curve close from about 1 percentage point to about half a percentage

Figure 13: job-finding rates conditional on unemployment

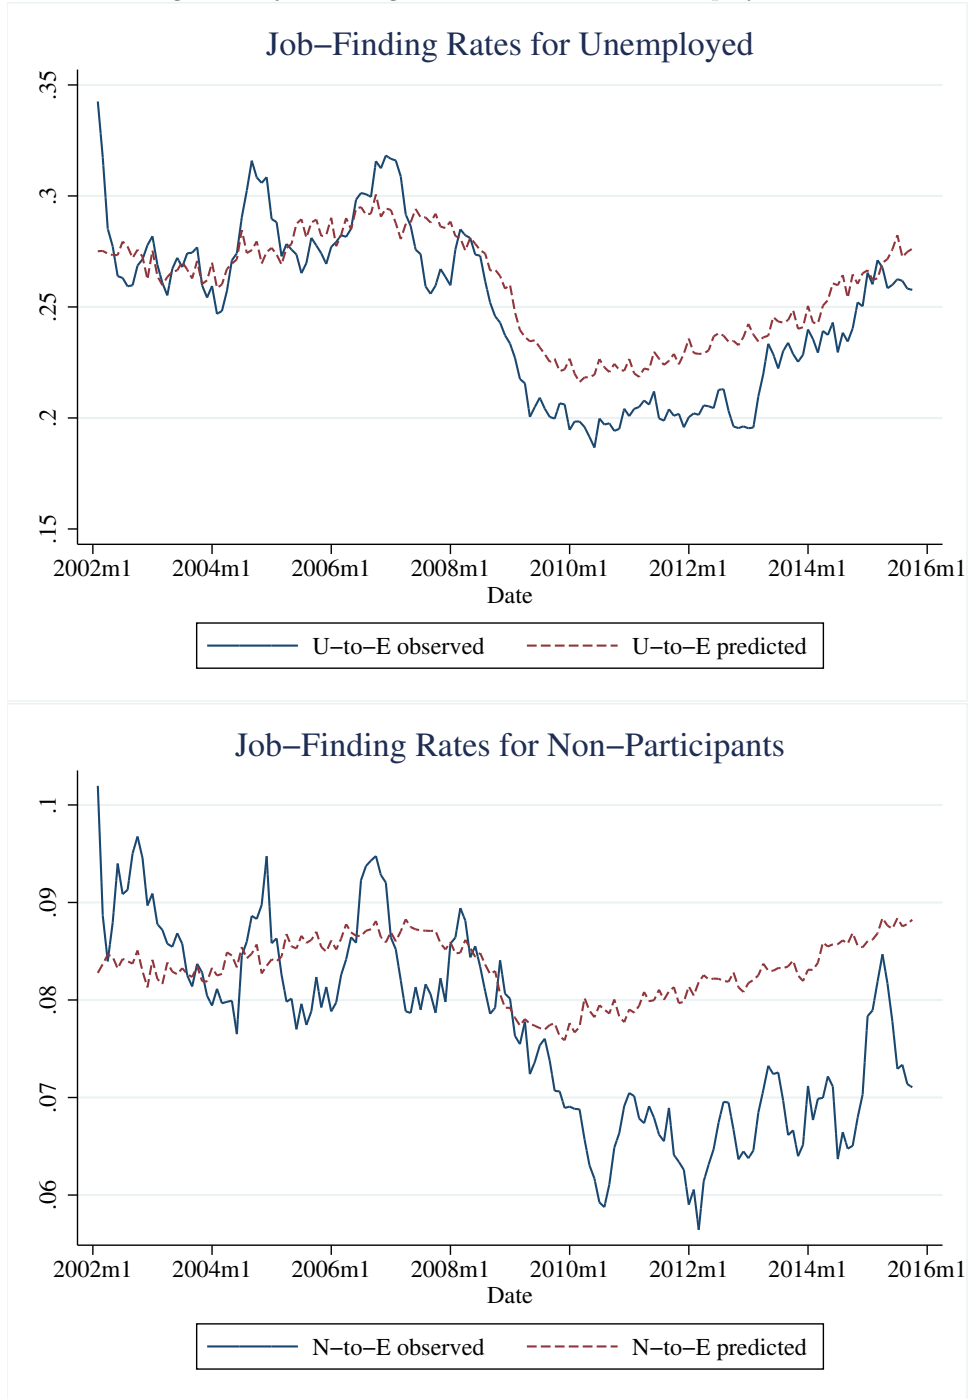
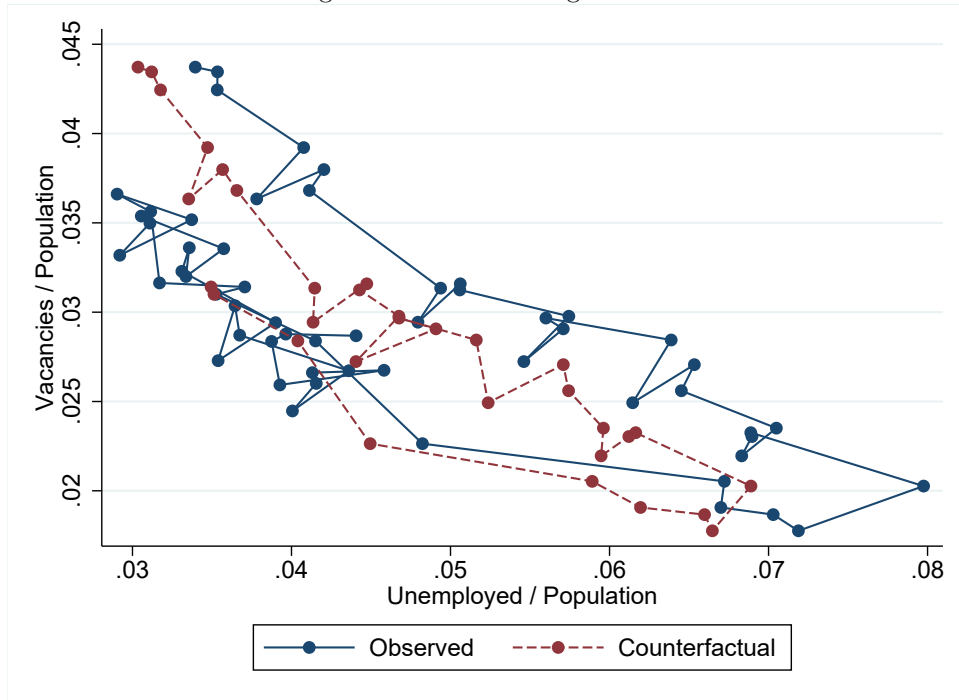


Figure 14: The Beveridge Curve

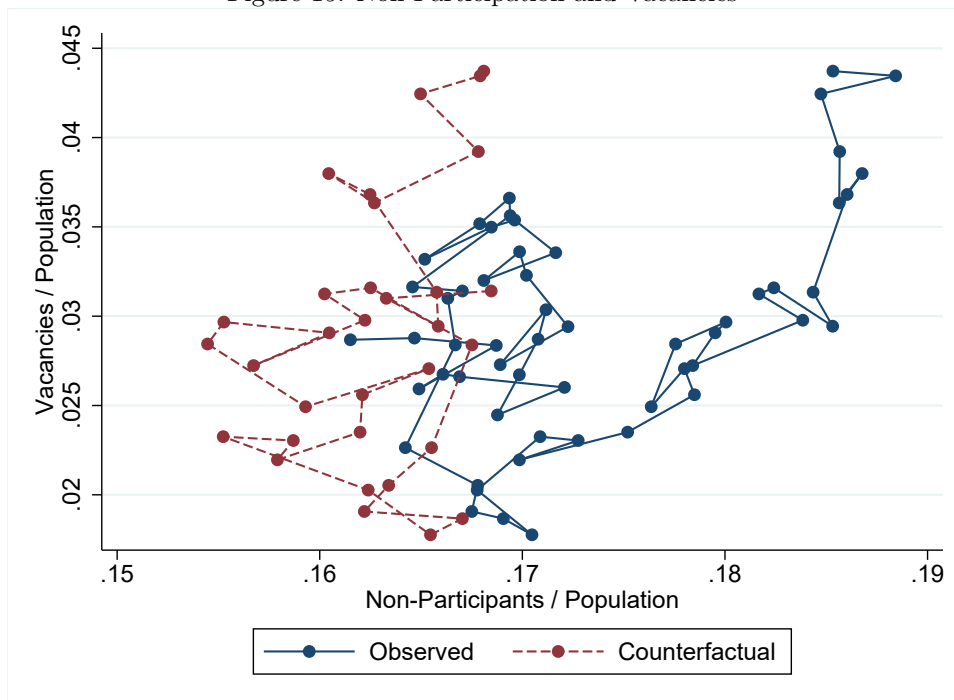


point. Overall, the model does a fairly good job of accounting for these additional 30 month of data on unemployment and vacancies.

As noted in Kroft et al. [2016], the model fails to explain by how much non-participation increased during the Great Recession. This failure of the model persists after April 2013 (see Figure 15). It should be noted however that the failure comes almost entirely from predicting a sharp decline in non-participation rather than the observed increase during the first half of the period up until about April 2013. Besides this shift early in the Great Recession, the model has tracked how the stock of non-participants evolved quite well.

Overall, we conclude that over the additional two years out-of-sample, the model calibrated on 2002-2007 data continues to perform very well when matching the dynamics in long-term unemployment, reasonably well when fitting the evolution of unemployment overall, but as before fails to match the persistent decline in the stock of non-participants.

Figure 15: Non-Participation and Vacancies



6 The Canadian Recession: A Laboratory to Understand Counterfactuals

In Kroft et al. [2016] (reviewed above), we augmented a standard matching model by allowing for non-participation as well as unemployment and by allowing for job-finding rates of the unemployed to exhibit duration dependence. We showed that this model was able to capture the increase in long-term unemployment over the Great Recession very well using variation in vacancies and job separation rates by type (into unemployment and non-employment) as the main ingredients. The model also captured about half of the shift in the Beveridge curve observed in the U.S. It however failed entirely to account for the increase in non-participation that characterized the U.S. Great Recession.

The CPS lacks data on the duration out of employment among the OLF. The model in Kroft et al. [2016] reflects this limitation of the CPS. The Canadian LFS however contains data on how long individuals have been without work, regardless of whether they are currently unemployed or OLF. This allows us to extend the model to account for a richer specification of how the duration of joblessness affects labor force transitions. The obvious extension is to allow the job-finding rate of both the unemployed and those OLF to depend on how long individuals have been out of employment. Thus, we model the job-finding rate for those OLF as follows:

$$\lambda^{NE}(d; x_t) = sB(d)m_0x_t^{1-\alpha} \quad (8)$$

where $B(d)$ captures the duration dependence in the job-finding rate among those OLF. This function is analogous to $A(d)$ in equation (2).

In addition, with the LFS, we can now model how transitions from between unemployment and OLF depend on duration of joblessness. To capture these, we model the duration-dependent transition rates from U to N ($\lambda_t^{UN}(d)$) and from N to U ($\lambda_t^{NU}(d)$) to be given by a time-specific shifter and a time-invariant function capturing the duration dependence. These duration-dependent transition rates are therefore

$$\lambda_t^{UN}(d) = \lambda_t^{UN}C(d) \quad (9)$$

and

$$\lambda_t^{NU}(d) = \lambda_t^{NU}D(d) \quad (10)$$

where $C(d)$ and $D(d)$ are functions capturing the duration structure in these transition rates. These functions are obtained by averaging the transition rates over the entire data period and then allowing for the relative transition rates to vary month-by-month up to duration 24. After 2 years of joblessness, transition rates are postulated not to vary with the length of joblessness any further.

6.1 Observed Duration Dependence in Labor Market Flows

Using the Canadian data between 1/2001 and 10/2015 in Canada, we can describe how joblessness durations affect labor market flows in a more encompassing manner, including flows originating among OLF. These flows are shown as functions of the duration of joblessness in Figure 16 for job-finding rates and Figure 17 for flows between U and N .

For both groups, the job-finding rate declines sharply over the first few months of joblessness and then flattens out somewhat. At a duration of one year, the job-finding rate is about half that of a newly jobless individual, whether the individual is unemployed or out of the labor force. We also observe strong duration patterns in flows between unemployment and non-participation. Figure 17 shows that the probability of transitioning from unemployment to non-participation increases and the probability of making a reverse transition from non-participation to unemployment declines steadily with joblessness duration. These findings are consistent with research in the U.S. that finds that individuals are increasingly likely to withdraw from the labor market the longer they are out of the job (Krueger et al. [2014]).

To capture these relationships, we estimate flexible nonlinear functions that capture the associations

Figure 16: job-finding Rate by Duration of Joblessness

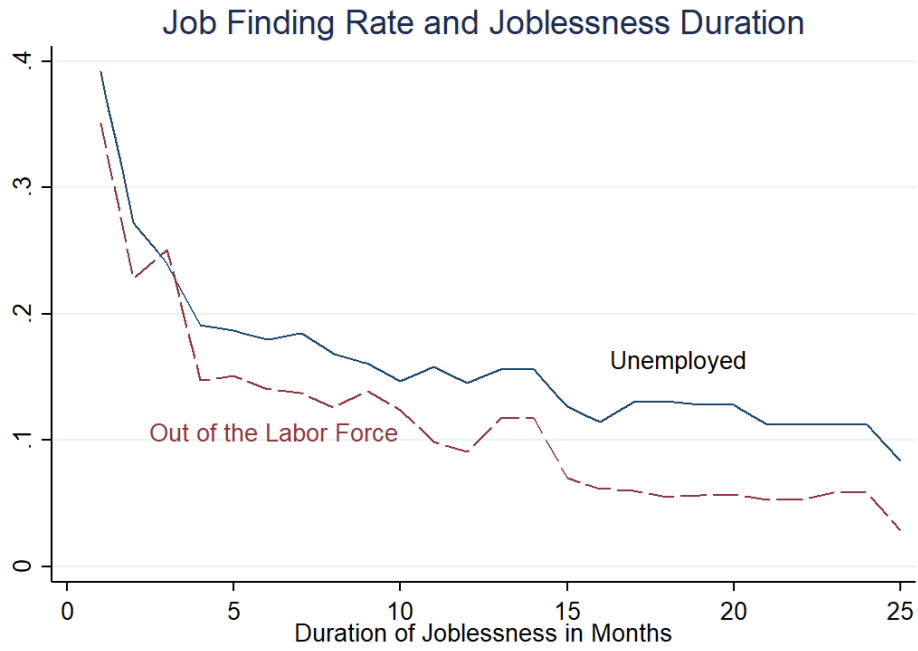


Figure 17: Transitions $U \rightleftharpoons N$ by Duration of Joblessness

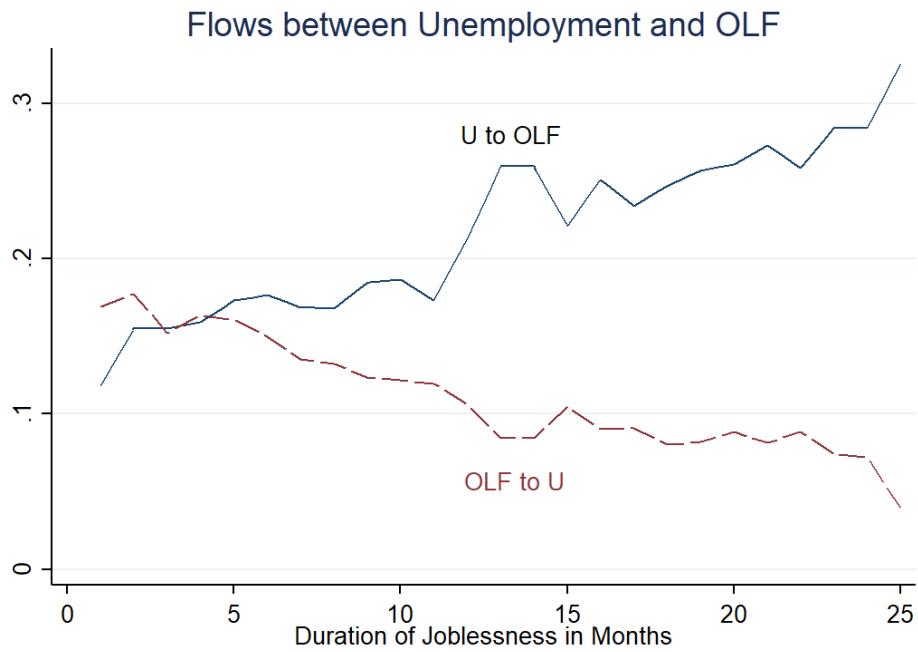


Table 1: Parameterizing the Matching Model with Duration Dependence in Joblessness

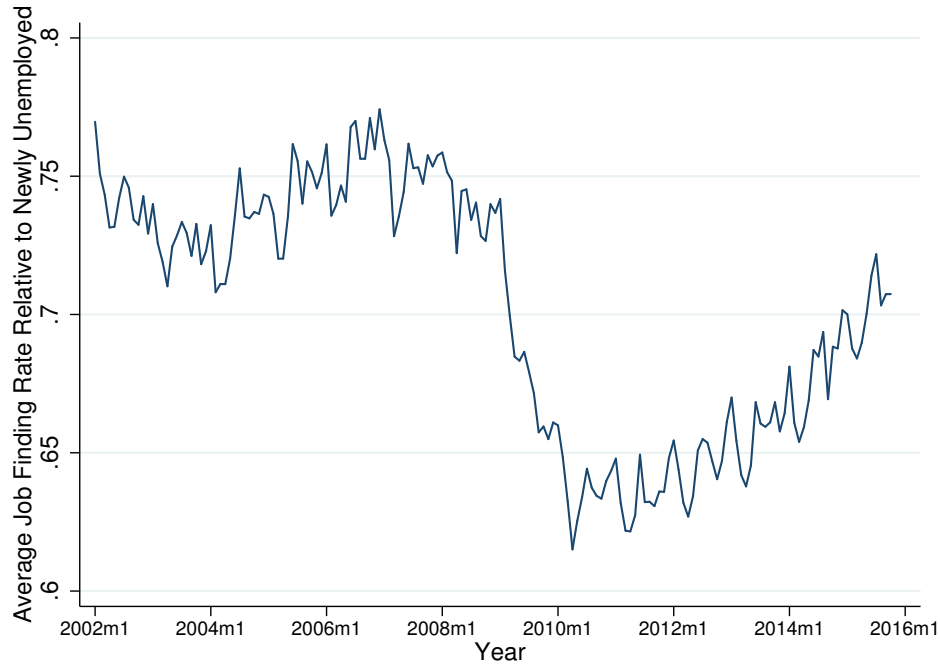
| Table 1 | | |
|---|------------------|---------------|
| Model-Based Estimates | | |
| Using Duration Jobless | | |
| Duration Dependence Parameters | | |
| a_1 (intercept parameter unemployed) | | 0.671 |
| a_1 (intercept parameter non participant) | | 0.862 |
| b_1 (slope parameter unemployed) | | 0.448 |
| b_1 (intercept parameter non participant) | | 0.283 |
| $A(d)$ or $B(d) = (1-a_1) + a_1 \cdot \exp(-b_1 \cdot D)$ | | |
| Matching Model Parameters | | |
| | No Duration in N | Duration in N |
| α | 0.406 | 0.430 |
| m_0 (scale parameter) | 0.128 | 0.200 |
| s (relative search intensity of inactive) | 0.175 | 0.723 |
| p | 7.558 | 7.404 |
| $M(U+sI, V) = m_0(U+sI)^\alpha V^{1-\alpha}$ | | |

Notes: This table reports estimates from model calibration to match the pre-recession data. These parameters are then used to construct counterfactual predictions of unemployment, employment, long-term joblessness, and other labor market outcomes. See text for details.

between joblessness duration and job-finding rates (for unemployment and non-participants). The specific two-parameter function follows Kroft et al. [2016] and is the following: $A(d) = (1 - a) + a \cdot \exp(-b \cdot d)$. These functions fit the job-finding rates across the duration distribution well (see Appendix Figure 37). The estimated parameters are reported in Table 1.

The parameters specified in $A(d)$ and $B(d)$ are estimated using non-linear least squares (NLLS) regression based on the “cell” averages (i.e., average job-finding rate by joblessness duration). The NLLS estimation is weighted using the number of observations at that duration. The functions are defined so that the function is defined relative to newly jobless individual (at $d = 0$). Both functions show steep declines (of more than 50 percent) over the first several months. For the unemployed, the function “flattens out” considerably after roughly six months, which mirrors what has been found in the U.S. (Kroft et al. [2016]). For non-participants, the function also steeply declines, but does not flatten out quite as rapidly, and also declines further in absolute terms: after 12 months, the job-finding rate of non-participants is less than 20 percent of a non-participant with a jobless duration of 0, while for unemployed

Figure 18: A Measure of the Search Relevant Duration Structure in Unemployment: $\overline{A(d)}$

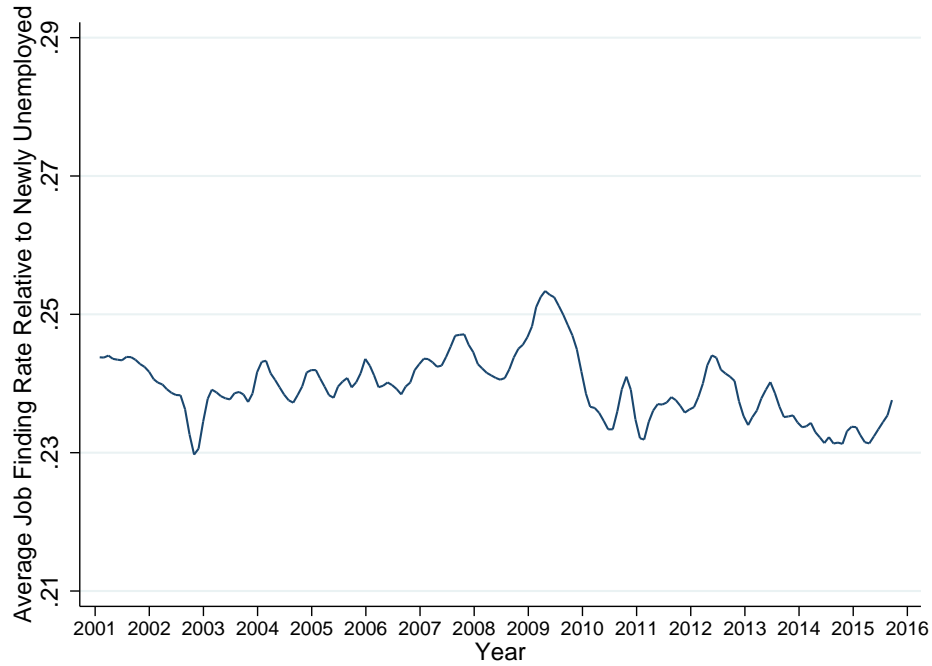


individuals the job-finding rate (at 12 months) is roughly 30 – 40 percent of the job-finding rate of an unemployed individual with duration of 0.

With these estimates, we can next examine how job-finding rates change over time based solely on changes to the distribution of joblessness. For the unemployed, this is the average of the $A(d)$ function (averaged across current stock of joblessness durations). Intuitively, in a recession, longer spells receive more weight and so pull down the mean $A(d)$. For Canada, the mean is roughly 52 percent in the pre-recession period and drops to about 48 percent during the recession as can be seen in Figure 18. For the U.S., the predicted job-finding rate is about 75 percent and falls to about 65 percent during the recession. Again, this is consistent with the trends in LTU/LTJ for both countries, showing that the recession was much deeper in the U.S. and generated much more long-term unemployment.

For the non-participants, we can likewise report how the job-finding rate conditional on non-participation evolved solely based on the distribution of the duration of joblessness. This measure is obtained by averaging $B(d)$ using the observed duration distribution among non-participants. Figure 19 shows that the average $B(d)$ is much lower than the average $A(d)$, starting around 23 percent in the pre-recession period and falling to about 21 – 22 percent during the recession. This reflects the fact that most of the jobless among the non-participants have very long spells and long-term joblessness increased only slightly for this group during the recession.

Figure 19: A Measure of the Search Relevant Duration Structure among OLF: $\overline{B(d)}$



6.2 Calibrating the model with Canadian data

To implement the matching model in equations (2) and (8) we require the parameters entering the matching function (s, p, m_0, α) as well as the functions $A(d)$, $B(d)$, $C(d)$, and $D(d)$. We described in the previous section the functional form and estimation of the functions $A(d)$ and $B(d)$. For the functions $C(d)$ and $D(d)$, we use the average transition rates in Figure 17 by month of joblessness duration without imposing a functional form. The functions $A(d)$, $B(d)$, $C(d)$, and $D(d)$ are estimated using the entire sample period from January 2001 to October 2015. To estimate the matching parameters (s, p, m_0, α) we only use the data up to the onset of the recession in October 2008. The parameter estimates of the matching model as well as of $A(d)$, $B(d)$ are shown in Table 1.

Figure 20 shows the observed and fitted job-finding rates conditional on unemployment from January 2001 to October 2015. Contrary to appearance, this figure shows not just the fitted value for the full model allowing for both duration dependence among the unemployed and those OLF, but it also shows the fit for a model that is restricted to allow only for duration dependence among the unemployed. The predicted values for these two models overlap almost perfectly. Overall, we see that the model fits the overall decline in job-finding rates quite well, with the exception of the initial decline in observed job-finding rates.

Figure 21 shows the observed and fitted job-finding rates among non-participants. Here we observe

Figure 20: Observed and Fitted job-finding Rates Conditional on Unemployment

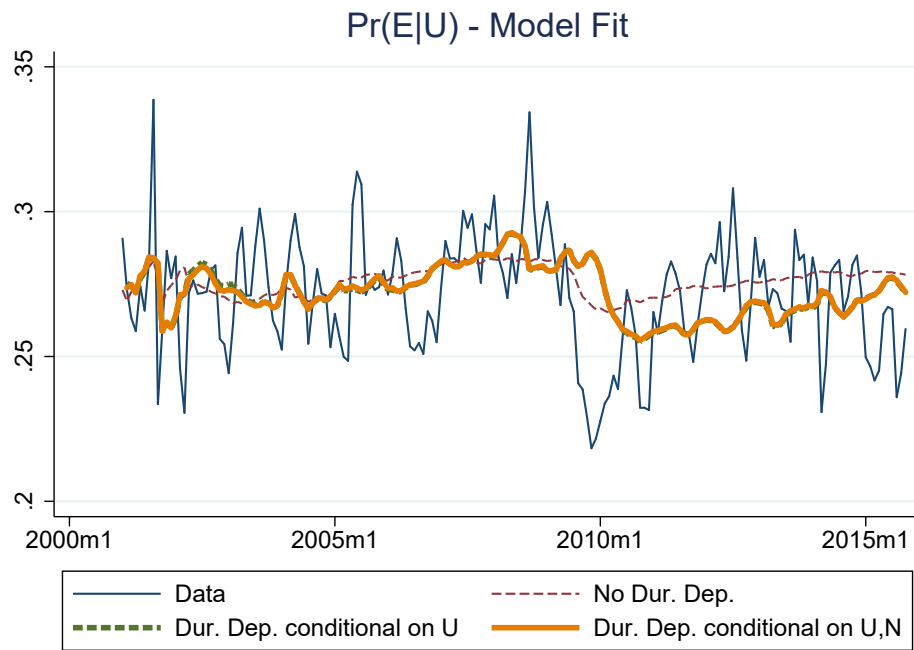
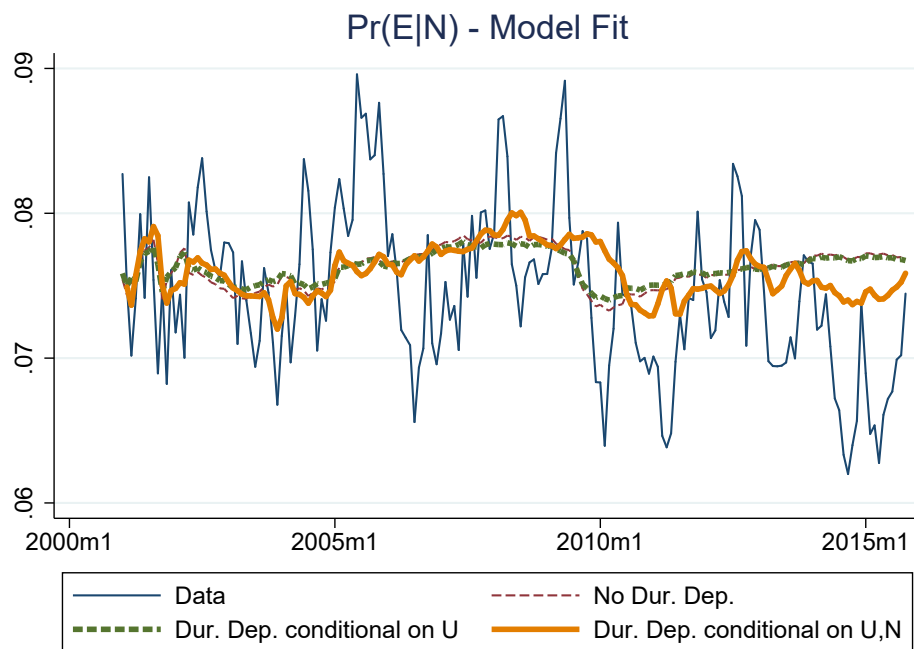


Figure 21: Observed and Fitted job-finding Rates Conditional on Non-participation



that allowing for duration dependence in both U and N improves the fit of the model. The thick line shows the fit of our preferred model, while the dashed line shows the fit without duration dependence among the OLF. The restricted model overestimates the job-finding rate among non-participants after 2012, while the unrestricted model does a reasonable job of fitting these later periods as well. Overall, we note that the model fits the decline in the job-finding rate quite well.

It is noteworthy that in Canada, but not in the U.S., the matching model estimated on the pre-recession period matches the job-finding rates even after the recession. This is true even if we estimate the restricted model in Kroft et al. [2016] rather than the model with enriched dynamics (eqs 8, 9, 10). In either case, the matching model fitted on Canadian data fits the job-finding rates both quantitatively and qualitatively, while in the U.S. we underestimate the decline in the job-finding rate especially among the OLF.

6.3 Counterfactual Results

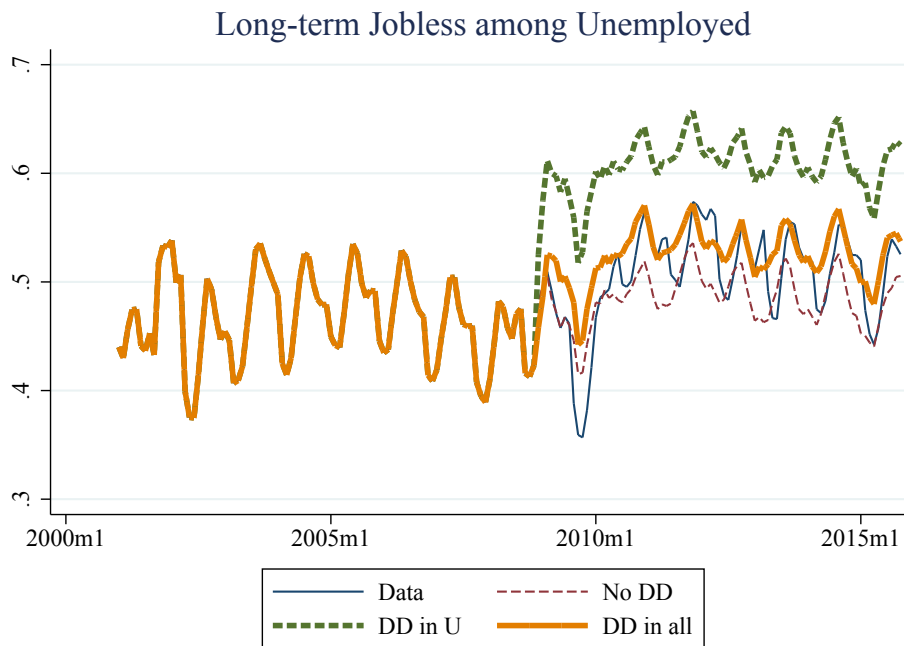
Note that Figures 20 and 21 were constructed using the observed distributions of the duration of joblessness across the entire time-period. These figures do therefore not speak to how well the model performs in generating counterfactual distributions across (U, N, E) and across the duration of joblessness. We now turn to counterfactual simulations that allow us to probe how well the model of labor force flows developed here describes the dynamics of the labor market overall.

In particular, we simulate the model using as exogenous forcing variables the vacancy series, the job-loss rates conditional on U and N , and the transition rates between U and N observed in the data.¹² We keep constant the parameters of the matching model (eqs. 2 and 8) and the relative transition rates $A(d)$, $B(d)$, $C(d)$, and $D(d)$.

We begin by comparing how the counterfactual pattern in long-term joblessness (>26 weeks) among the unemployed and those OLF evolves. Figure 22 shows the share of long-term joblessness (>26 weeks) among unemployed in the data as well as for 3 different counterfactual model specifications after 10/2008. The dashed thin line refers to the model without any duration dependence. Surprisingly, this model does fairly well in fitting the long-term joblessness distribution conditional on unemployment. The thick dashed line (labeled “DD in U”) shows the counterfactuals based on a model that allows only for duration dependence among the job-finding rates of the unemployed, but not among non-participants, nor in $U \rightleftharpoons N$ flows. This model predicts a much larger share of long-term joblessness than observed in the data. It does so, because it does not account for the fact that many of those unemployed for long periods

¹²For the latter, we take the average transition rates observed in the data $P(U|N)$ and $P(N|U)$ as well as the duration distribution in the data to construct the rates of transition between U and N conditional on duration 0. These are used as the forcing variables in the dynamic simulation.

Figure 22: Long-term Joblessness among Unemployed in Canada (>26 weeks)



become non-participants (see Figure 17). By contrast the full model allowing for duration dependence in job-finding rates among both U and OLF and also for duration dependence in flows $U \rightleftharpoons N$ fits the data very well.

Figure 23 the same data and counterfactuals for the long-term jobless conditional on being OLF . We observe that models that do not allow for duration dependence in the job-finding rates of N as well as in $U \rightleftharpoons N$ underestimate the share of long-term joblessness among the non-participants. This is so, because they do not capture the large unbalanced flows between unemployment and non-participation at high durations of joblessness. Once we allow for duration dependence in these flows, we can fit the patterns in long-term joblessness quite well using the stable matching framework and a duration structure that is unchanged over the study period.

Another way of assessing the performance of the model present the counterfactual stocks of employment and unemployment during the recession. These are shown in Figures 24 and 25. Here we see that all our models trace out a path for employment that exceeds the observed path and a path for unemployment that fall short of that actually observed. Part of this comes from the fact that the model-based job-finding rate was predicted to be higher than was actually the case during the recession.

Figure 26 traces out the counterfactual Beveridge curve in unemployment using the full model with duration dependence in all flows since the onset of the recession in October 2008. Also displayed is

Figure 23: Long-Term Joblessness among OLF in Canada (>26 weeks)

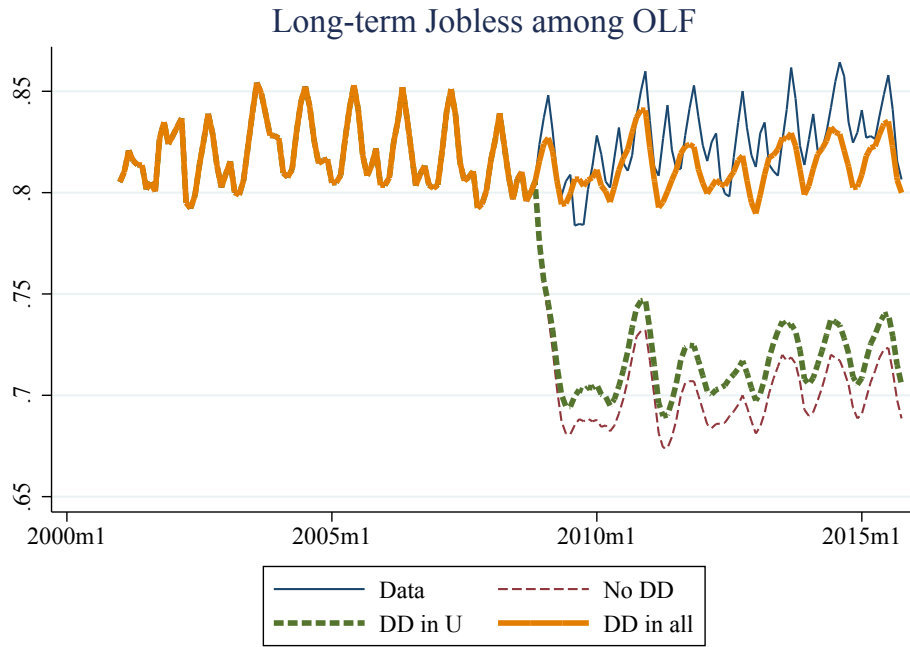


Figure 24: Employment during the Canadian Recession

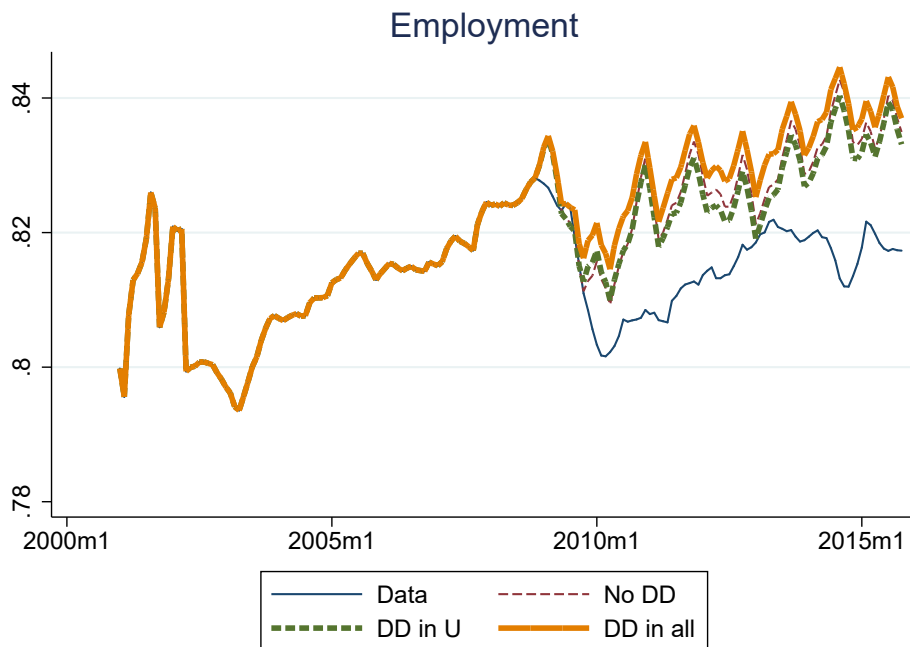
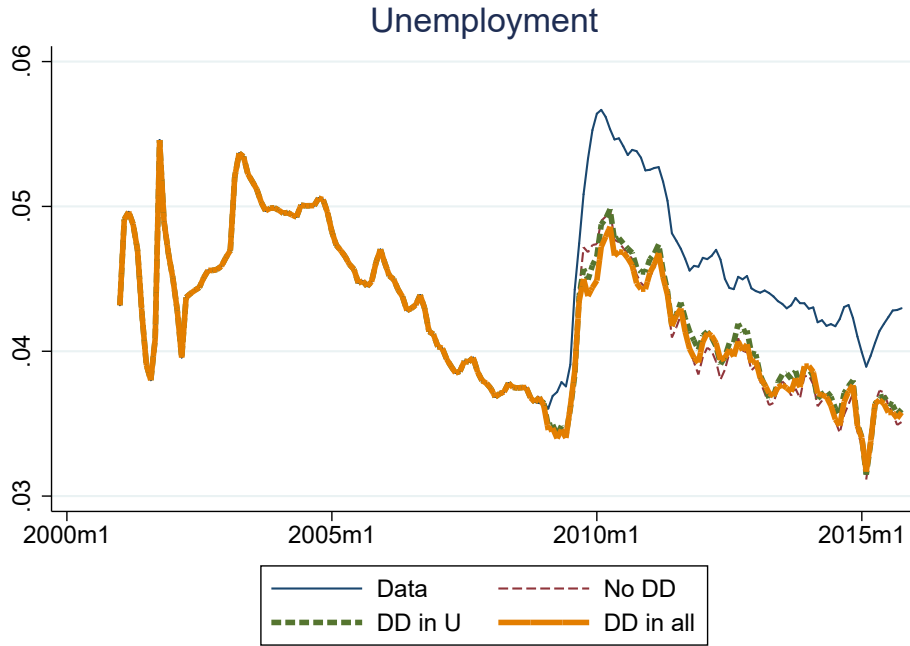


Figure 25: Unemployment during the Canadian Recession



the observed Beveridge Curve over this time period. We see that the model fits the data quite closely. Overall, we observe that a difference of less than half a percentage point in the unemployment-population rate over the Great recession. And, the counterfactual series traces out a similarly sloped relationship between unemployment and vacancies over this time-period. The counterfactual performance of the full model on the Canadian data is thus quite a bit superior to the performance of the U.S. model in Kroft et al. [2016]. The fit of the counterfactual N-V curve (see Figure 27) is not as good as the fit of the Beveridge curve in unemployment in that the model persistently underestimates the share OLF by about 1 percentage point. However, we again observe that the fit is superior to the fit in the U.S. model in Kroft et al. [2016].

Thus, we have so far seen that the dynamics of long-term joblessness conditional on unemployment and OLF are very sensitive to correctly modeling the duration dependence in both job-finding rates and in the flows between U and OLF. Overall, the full model does also quite well in matching the stocks in U and OLF over the recession.

6.4 A Canadian Great Recession in Labor Demand

The final question we want to take up using our full counterfactual model is whether the Canadian labor market performed better than the U.S. labor market primarily because of differences in the demand for

Figure 26: A Counterfactual Beveridge Curve - Canada

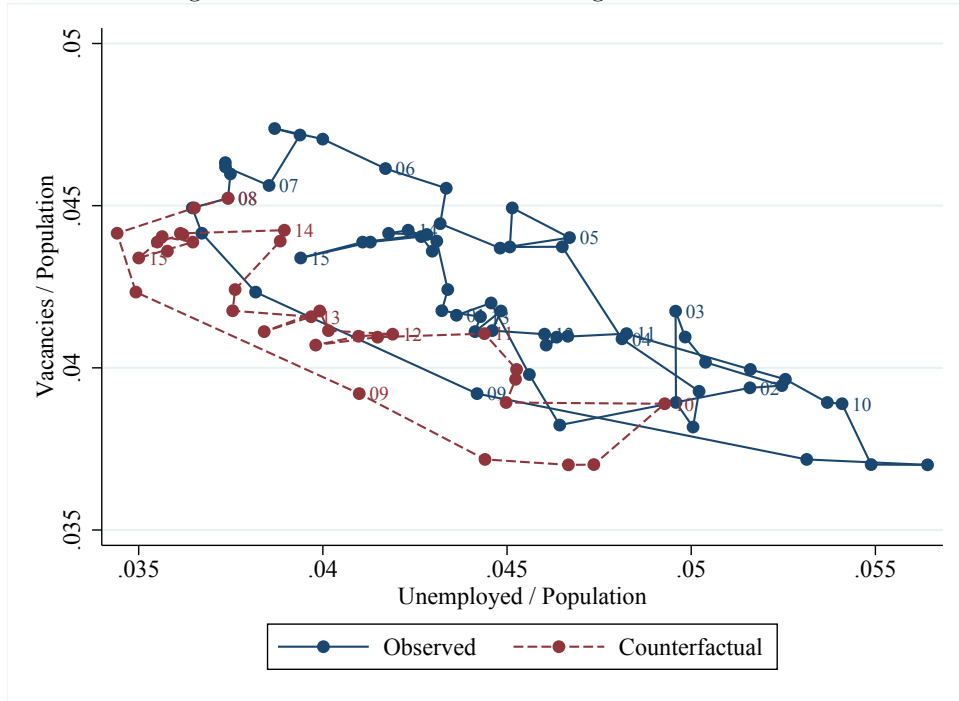
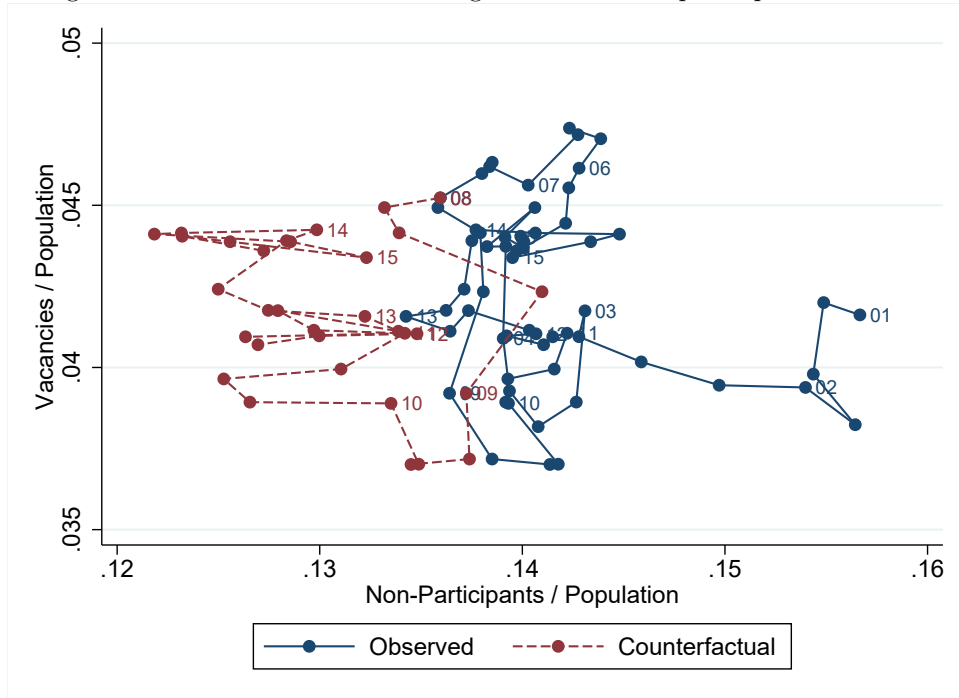


Figure 27: A Counterfactual Beveridge Curve for Non-participants - Canada



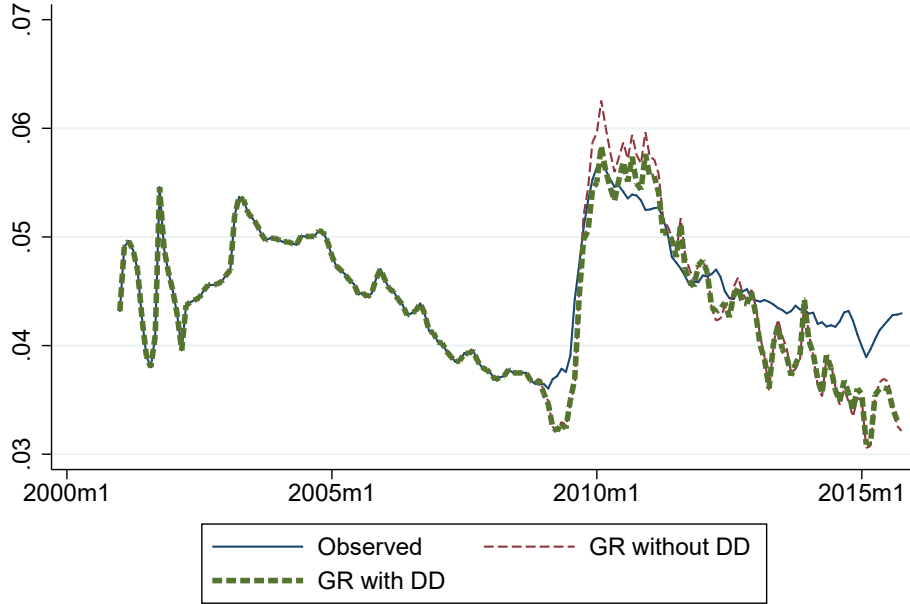
labor over the recessionary period. That is, maybe Canada avoided a Great recession because Canadian employers continued to create vacancies at higher rates and because Canadian employers did not lay-off quite as many workers during the recession itself. To investigate this question, we take the counterfactual model that we saw fits the data quite well over the recession and apply to it the same time-series pattern in job separation rates and in vacancies observed in the U.S. during the Great recession. That is, we apply the percentage decrease in vacancies observed during the Great recession to the vacancy series in Canada starting in October 2008 and we likewise apply the increase in the job loss rate over the Great recession. To illustrate, the vacancy rate in the U.S. in December 2008, one year after the onset of the Great recession stood at 66% of the vacancy rate in December 2007. In our next set of counterfactual simulations we thus postulate that the vacancy rate in Canada in Oct. 2009 stands at 66% of that in Oct. 2008, the onset of the Canadian recession. At that point, the vacancy rate in Canada had in fact only declined by about 19%. We proceed in the same fashion for job loss rates and thus impose higher job losses on the Canadian economy as well as more long-lasting job losses.

The two panels in Figure 28 show the implied time-series of unemployment and non-participation from the full counterfactual model as well as the counterfactual model that does not allow for duration dependence. By coincidence, the counterfactual models in fact reproduce the observed patterns in the unemployment rate in Canada, but that finding is without significance. More meaningful are two observations. The first observation is that even under the conditions of the Great recession in terms of labor demand, the full model and the model without any duration dependence do not diverge widely in predicted unemployment and OLF patterns over the recession. This reinforces the conclusion that duration dependence is significantly more important for understanding the patterns in long-term joblessness than it is for understanding the overall performance of the labor market in terms of the broad labor force stocks (E,U,N). We believe this to be the case because these latter stocks and their dynamics are quite sensitive to the flows early during jobless spells, since the majority of jobless spells are relatively short.

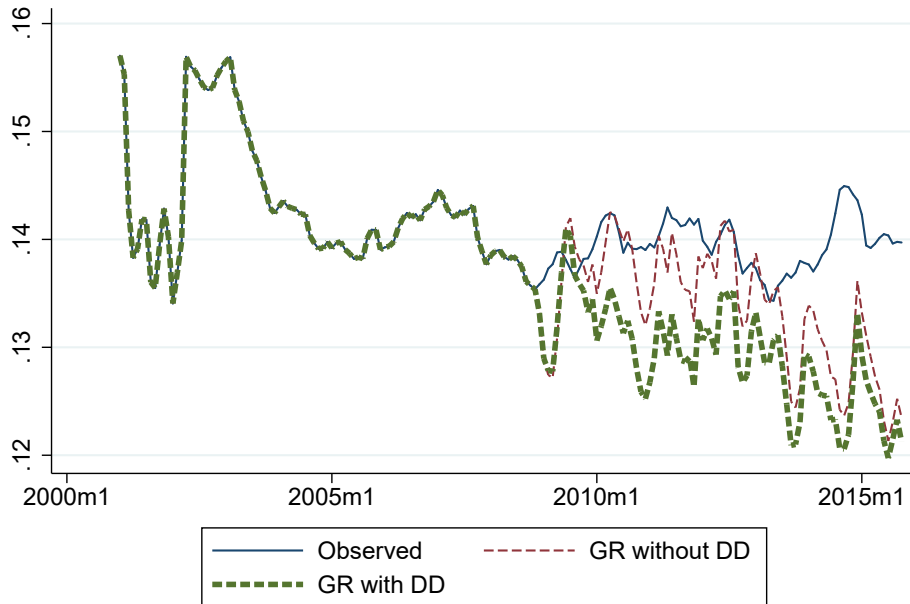
The second observation based on figure 28 is that the much greater decline in labor demand and the larger and more persistent increase in job loss rates in the U.S. compared to Canada is not enough to explain the different labor market experiences of the two countries over the last decade. Clearly, imposing U.S. style deterioration in labor demand would have significantly worsened the recession in Canada, but it does not seem as if this would have induced a Great Canadian recession. Rather, it seems that the recession in Canada was mild compared to the recession in the U.S. in large part because the rate at which job seekers were matched to vacancies did not break down to the same extent in Canada as it did in the U.S. The matching model was more stable in Canada.

Figure 28: A Canadian Great Recession

A Canadian Great Recession: Unemployment



A Canadian Great Recession: OLF



7 Conclusion

This paper showed that long-term unemployment rose sharply in both Canada and the U.S. during the Great Recession. The levels of long-term unemployment continue to remain elevated in both countries and do not appear to be explained by shifts in the observable characteristics of the unemployed.

Thus, similar to our previous analysis of the Great Recession in the U.S., we turned to a matching model to try to understand the recent labor market dynamics of the unemployed. The Canadian labor market data permitted us to extend the matching model in one important direction. In particular, since we observe joblessness spells for both unemployed individuals and non-participants, we can model duration dependence in the job-finding rate for both groups, something we were unable to do with the U.S. analysis. Additionally, we can allow for duration dependence in all flows between unemployment and non-participation. Our results indicate strong negative duration dependence in job-finding rates for non-participants and also strong duration dependence in flows between unemployment and non-participation. When we broaden the matching model to include these features for Canada, we find that the fit of the model improves substantially. Overall, we interpret these results as providing evidence that joblessness is a better measure of non-employment when analyzing labor market trends. Finally, a contribution of the paper has been to introduce a new measure of vacancies which we hope will spur other researchers in Canada to consider questions relating to the role of labor demand in Canada.

References

- Daniel Aaronson, Jonathan Davis, and Luojia Hu. Explaining the decline in the u.s. labor force participation rate. *Federal Reserve Bank of Chicago Chicago Fed Letter*, 296, 2012.
- Hie Joo Ahn and James D Hamilton. Heterogeneity and unemployment dynamics. Technical report, National Bureau of Economic Research, 2016.
- Joseph G. Altonji, Prashant Bharadwaj, and Fabian Lange. Changes in the characteristics of american youth: Implications for adult outcomes. *Journal of Labor Economics*, 30(4):783 – 828, 2012.
- Patricia M Anderson and Simon M Burgess. Empirical matching functions: estimation and interpretation using state-level data. *Review of Economics and Statistics*, 82(1):93–102, 2000.
- David Autor and Mark Duggan. The growth in social security disability rolls: A fiscal crisis unfolding. *Journal of Economic Perspectives*, 20(3):71–96, 2006.
- David Autor, Nicole Maestas, Kathleen Mullen, and Alexander Strand. Does delay cause decay? the effect of administrative decision time on the labor force participation and earnings of disability applicants. *Working Paper*, 2013.
- Regis Barnichon. Vacancy posting, job separation, and unemployment fluctuations. *Journal of Economic Dynamics and Control*, 36(3):315–330, 2012.
- Regis Barnichon and Andrew Figura. Labor market heterogeneity and the aggregate matching function. *Working Paper*, 2013.
- Andrew Barr and Sarah Turner. Down and enrolled: An examination of the enrollment response to cyclical trends and job loss. *unpublished Manuscript, University of Virginia*, 2012.
- Una-Louise Bell et al. A comparative analysis of the aggregate matching process in france, great britain and spain. Technical report, Banco de España, 1997.
- Leila Bengali, Mary Daly, and Rob Valletta. Will labor force participation bounce back? *Federal Reserve Bank of San Francisco Economic Letter*, May 2013.
- Samuel Bentolila, Ignacio García-Pérez, and Marcel Jansen. Are the spanish long-term unemployed unemployable. *SERIE*, 8, 2017.
- Andre Bernard and Jeannine Usalca. The labour market in canada and the us since the last recession. *Economic Insights*, (36), 2014.
- Olivier J. Blanchard and Peter Diamond. Ranking, unemployment duration, and wages. *Review of Economic Studies*, 61(3):417–434, 1994.
- Olivier Jean Blanchard, Peter Diamond, Robert E Hall, and Janet Yellen. The beveridge curve. *Brookings papers on economic activity*, 1989(1):1–76, 1989.
- John E. Bregger and Steven E. Haugen. Bls introduces new range of alternative unemployment measures. *Monthly Labor Review*, 118:19–26, 1995.
- Simon M Burgess. A model of competition between unemployed and employed job searchers: an application to the unemployment outflow rate in britain. *The Economic Journal*, 103(420):1190–1204, 1993.
- Maria Canon, Marianna Kudlyak, and Marisa Reed. Is involuntary part-time employment different after the great recession? *The Regional Economist, Federal Reserve Bank of St. Louis*, July 2014.
- David Card and Craig W. Riddell. A comparative analysis of unemployment in canada and the united states. *NBER Chapters, in: Small Differences That Matter: Labor Markets and Income Maintenance in Canada and the United States*, pages 149–190, 1993.

- David Card, J. Kluge, and A. Weber. Active labour market policy evaluations: A meta-analysis. *The Economic Journal*, 120:452–477, November 2010.
- Kim B. Clark and Lawrence H. Summers. Labor market dynamics and unemployment: A reconsideration. *Brookings Papers on Economic Activity*, 1:13–60, 1979.
- Bruno Crepon, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora. Do labor market policies have displacement effect? evidence from a clustered randomized experiment. *The Quarterly Journal of Economics*, 128(2):531–580, 2013.
- Steven J. Davis, Jason R. Faberman, and John C. Haltiwanger. The establishment-level behavior of vacancies and hiring. *Quarterly Journal of Economics*, 128(2):581–622., 2013.
- Peter Diamond. Cyclical unemployment, structural unemployment. *NBER Working Paper*, 18761, 2013.
- Michael W. Elsby, Ryan Michaels, and Gary Solon. The ins and outs of cyclical unemployment. *American Economic Journal: Macroeconomics*, 1(1):84–110, 2009.
- Michael W. Elsby, Bart Hobijn, and Aysegul Sahin. The labor market in the great recession. *Brookings Papers on Economic Activity*, pages 1–69, 2010.
- Michael W. Elsby, Bart Hobijn, Aysegul Sahin, and Robert G. Valletta. The labor market in the great recession - an update to september 2011. *Brookings Papers on Economic Activity*, pages 353–371, 2011.
- Michael W. Elsby, Bart Hobijn, and Aysegul Sahin. On the importance of the participation margin for labor market fluctuations. *Federal Reserve Bank of San Francisco Working Paper Series*, 2013-05, 05 2013a.
- Michael W. Elsby, Bart Hobijn, and Aysegul Sahin. Unemployment dynamics in the oecd. *Review of Economics and Statistics*, 95(2):530–548, 2013b.
- Christopher Erceg and Andrew Levin. Labor force participation and monetary policy in the wake of the great recession. *Unpublished manuscript*, April 2013.
- Stefan Eriksson and Dan-Olof Rooth. Do employers use unemployment as a sorting criterion when hiring? evidence from a field experiment. *American Economic Review*, forthcoming.
- Henry S. Farber and Robert G. Valletta. Do extended unemployment benefits lengthen unemployment spells? evidence from recent cycles in the u.s. labor market. *Federal Reserve Working Paper*, 2013.
- Henry S Farber, Dan Silverman, and Till Von Wachter. Factors determining callbacks to job applications by the unemployed: An audit study. Technical report, National Bureau of Economic Research, 2015.
- Christopher J. Flinn and James J. Heckman. Are unemployment and out of the labor force behaviorally distinct labor force states? *Journal of Labor Economics*, 1(1):28–42, 1983.
- Harley J. Frazis, Edwin L. Robinson, Thomas D. Evans, and Martha A. Duff. Estimating gross flows consistent with stocks in the cps. *Monthly Labor Review*, 128(9):3–9, 2005.
- Shigeru Fujita. Constructing the reason-for-nonparticipation variable using the monthly cps. 2014a.
- Shigeru Fujita. On the causes of declines in the labor force participation rate. *Federal Reserve Bank of Philadelphia*, 2014b.
- Shigeru Fujita and Giuseppe Moscarini. Recall and unemployment. 2013.
- Shigeru Fujita and Garey Ramey. The cyclical of separation and job finding rates. *International Economic Review*, 50(2):414–430, 2009.
- Rand Ghayad. Why is it so hard for job-seekers with valuable experience to find work? evidence from a field experiment. *Northeastern Working Paper*, 2013.

- Rand Ghayad and William Dickens. What can we learn by disaggregating the unemployment-vacancy relationship? *Federal Reserve Working Paper*, 2012.
- Robert J. Gordon. The phillips curve is alive and well: Inflation and the nairu during the slow recovery. *NBER Working*, (19390), 2013.
- Reuben Gronau. Wage comparison - a selectivity bias. *Journal of Political Economy*, 82(6):1119–1141, 1974. First paper to point out that job search strategies can induce variation in observed wage distributions from wage offer distributions.
- Robert E Hall. Job loss, job finding, and unemployment in the u.s. economy over the past fifty years. *NBER Macroeconomics Annual*, 20:101–137, 2005.
- Robert E. Hall and Sam Schulhofer-Wohl. Measuring matching efficiency with heterogeneous jobseekers. *mimeo*, 2013.
- James D Hamilton. Time series analysis. princeton university press: Princeton nj. *Harvey, A.C. Time Series Models. Halsted Press: New York NY*, 1981, 1994.
- James J. Heckman and Burton Singer. The identifiability of the proportional hazard model. *Review of Economic Studies*, 51(2):231–241, 1984.
- Richard Hendry, James A. Riccio, Richard Dorsett, David Greenberg, Genevieve Knight, Joan Phillips, Philip K. Robins, Sandra Vegeris, and Johanna Walter. Breaking the low-pay, no-pay cycle: Final evidence from the uk employment retention and advancement (era) demonstration. Technical Report 765, Department for Work and Pensions, UK, 2011. Experiment on job search assistance and incentive for work programs. Low effectiveness on single parents. Effective on LTU.
- Andreas Hornstein. Accounting for unemployment: The long and short of it. *Reserve Bank of Richmond Working Paper*, 12(07), 2012.
- Andreas Hornstein and Marianna Kudlyak. Estimating matching efficiency with variable search effort. Technical report, Mimeo, 2015.
- Andreas Hornstein, Per Krusell, and Giovanni L Violante. Frictional wage dispersion in search models: A quantitative assessment. Technical report, National Bureau of Economic Research, 2007.
- Andreas Hornstein, Marianna Kudlyak, and Fabian Lange. A new measure of resource utilization in the labor market. *Federal Reserve Bank of Richmond, mimeo*, April 2014a.
- Andreas Hornstein, Marianna Kudlyak, and Fabian Lange. Measuring resource utilization in the labor market. *Economic Quarterly*, 100(1):1 – 21, 2014b.
- Andreas Hornstein, Marianna Kudlyak, Fabian Lange, and Tim Sablik. Does the unemployment rate really overstate labor market recovery? *Federal Reserve Bank of Richmond Economic Brief*, June 2014c.
- Julie L. Hotchkiss. Changes in the aggregate labor force participation rate. *Federal Reserve Bank of Atlanta Economic Review*, 4:1–6, 2009.
- Gregor Jarosch and Laura Pilossoph. Statistical discrimination and duration dependence in the job finding rate. 2016.
- Stephen R. G. Jones and Craig W. Riddell. The measurement of unemployment: An empirical approach. *Econometrica*, 67(1):147–162, 1999.
- Stephen R. G. Jones and Craig W. Riddell. Unemployment and nonemployment: Heterogeneities in labor market states. *The Review of Economics and Statistics*, 88(2):314–323, 2006.
- Michael Joyce, Jerry Jones, and Jonathan Thomas. Non-employment and labour availability. *Bank of England Quarterly Bulletin*, 2003.

- Lawrence F. Katz. Layoffs, recall and the duration of unemployment. *NBER WP*, 1825, 1986.
- Lawrence F. Katz and Bruce D. Meyer. Unemployment insurance, recall expectations, and unemployment outcomes. *Quarterly Journal of Economics*, 105(4):973–1002, 1990.
- Lawrence F. Katz and Kevin M. Murphy. Changes in relative wages, 1963-1987: Supply and demand factors. *Quarterly Journal of Economics*, 107(1):35–78, 1992.
- Jochen Kluge. The effectiveness of european active labor market programs. *Labour Economics*, 17: 904–918, 2010.
- Kory Kroft, Fabian Lange, and Matthew J. Notowidigdo. Duration dependence and labor market conditions: Evidence from a field experiment. *Quarterly Journal of Economics*, 128(3):1123–1167, 2013.
- Kory Kroft, Fabian Lange, Matthew J. Notowidigdo, and Lawrence F. Katz. Long-term unemployment and the great recession: The role of composition, duration dependence, and non-participation. *Journal of Labor Economics*, 34(S1):S7–S54, 2016.
- Alan B. Krueger, Judd Cramer, and David Cho. Are the long-term unemployed on the margins of the labor market? *Brookings Papers on Economic Activity*, Spring 2014.
- Marianna Kudlyak and Fabian Lange. Measuring heterogeneity in job finding rates among those out of the labor force. *mimeo*, 2017.
- Rafael Lalive, Camille Landais, and Josef Zweimüller. Market externalities of large unemployment insurance programs. *IZA Working Paper*, (7650), 2013. Large Expansion of UI benefit in Austria - effects on non-eligible and eligible. Finds large effects on non-eligible workers.
- Camille Landais, Pascal Michaillat, and Emmanuel Saez. A macroeconomic approach to optimal unemployment insurance ii: Applications. *Working Paper*, 2016.
- Brigitte C. Madrian and Lars John Lefgren. A note on longitudinally matching current population survey (cps) respondents. *NBER Working Paper*, (T0247), 1999.
- Dale T. Mortensen and Christopher A. Pissarides. Job creation and job destruction in the theory of unemployment. *Review of Economic Studies*, 61(3):397–415, 1994.
- Arash Nekoei and Andrea Weber. Does extending unemployment benefits improve job quality? *American Economic Review*, 107(2):527 – 561, 2017.
- Congressional Budget Office. Understanding and responding to persistently high unemployment. 2012.
- Arthur M. Okun. Upward mobility in a high-pressure economy. *Brookings Papers on Economic Activity*, 4(1):207–262., 1973.
- Raymond E. Owens and Pierre-Daniel G. Sarte. How well do diffusion indexes capture business cycles? *Federal Reserve Bank of Richmond Economic Quarterly*, 91(4):23–42, 2005.
- Barbara Petrongolo and Christopher A Pissarides. Looking into the black box: A survey of the matching function. *Journal of Economic literature*, 39(2):390–431, 2001.
- Christopher Pissarides. Unemployment and vacancies in britain. *Economic policy*, 1(3):499–541, 1986.
- Christopher A. Pissarides. Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *American Economic Review*, 75:676–90, 1985.
- Anne E. Polivka and Stephen M. Miller. The cps after the redesign: Refocusing the economic lens. *Labor Statistics Measurement Issues; Studies in Income and Wealth*, 60, 1998.
- Layard Richard, Stephen Nickell, and Richard Jackman. Unemployment: Macroeconomic performance and the labour market, 1991.

- Jesse Rothstein. Unemployment insurance and job search in the great recession. *Brookings Papers on Economic Activity*, 43(2):143–213., 2011.
- Aysegul Sahin, Joseph Song, Giorgio Topa, and Giovanni L. Violante. Mismatch unemployment. *American Economic Review*, forthcoming.
- Johannes Schmieder, Till von Wachter, and Stefan Bender. The effect of unemployment benefits and nonemployment durations on wages. *American Economic Review*, 106(3):739 – 777, 2016.
- Mark Schweitzer. Ready, willing, and able? measuring labour availability in the uk. *Federal Reserve Bank of Cleveland Working Paper*, 0303, 2003.
- Robert Shimer. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49., 2005.
- Robert Shimer. Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15: 127–148, 2012.
- JanC van Ours. The efficiency of the dutch labour market in matching unemployment and vacancies. *De Economist*, 139(3):358–378, 1991.
- Marcelo Veracierto. Worker flows and matching efficiency. *Federal Reserve Bank of Chicago Economic Perspectives*, 35(4):147–169, 2011.
- Mark W Watson. Inflation persistence, the nairu, and the great recession. *The American Economic Review*, 104(5):31–36, 2014.

8 Appendix Figures

Figure 29: Long-Term Unemployment by Education

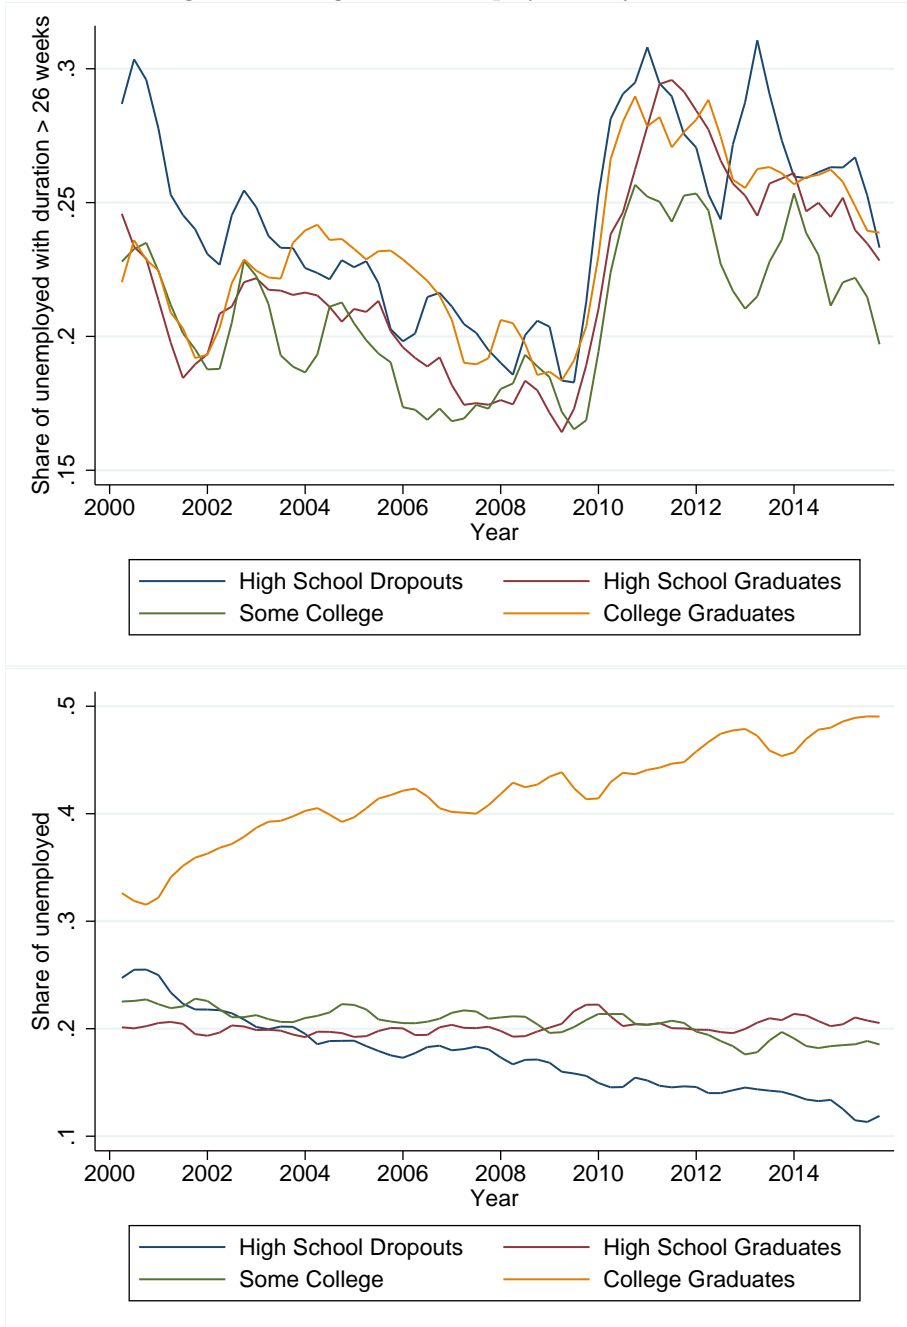


Figure 30: Long-Term Unemployment by Gender

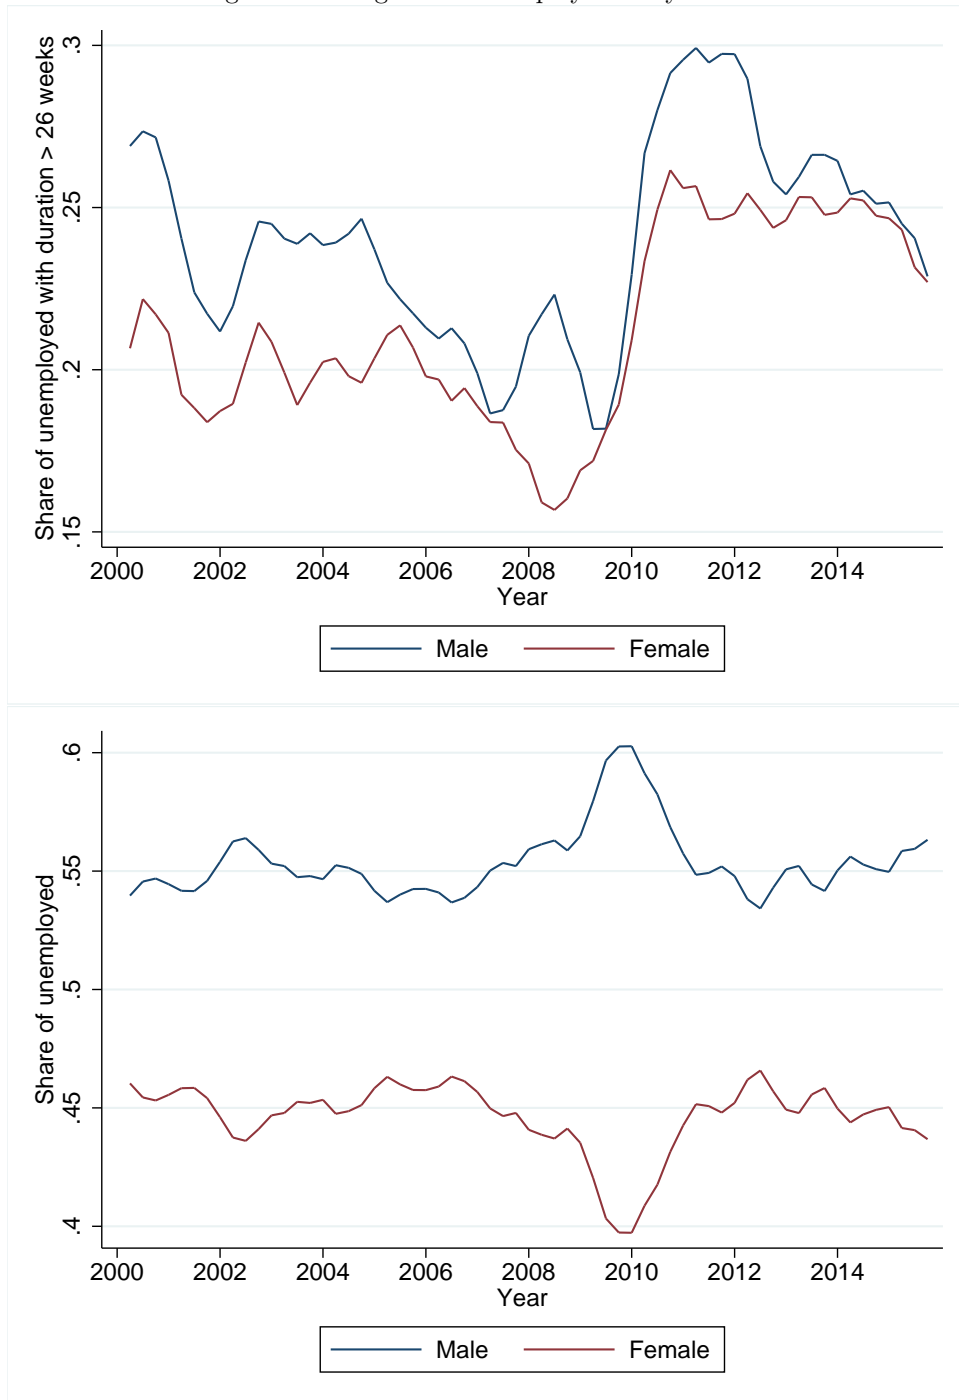
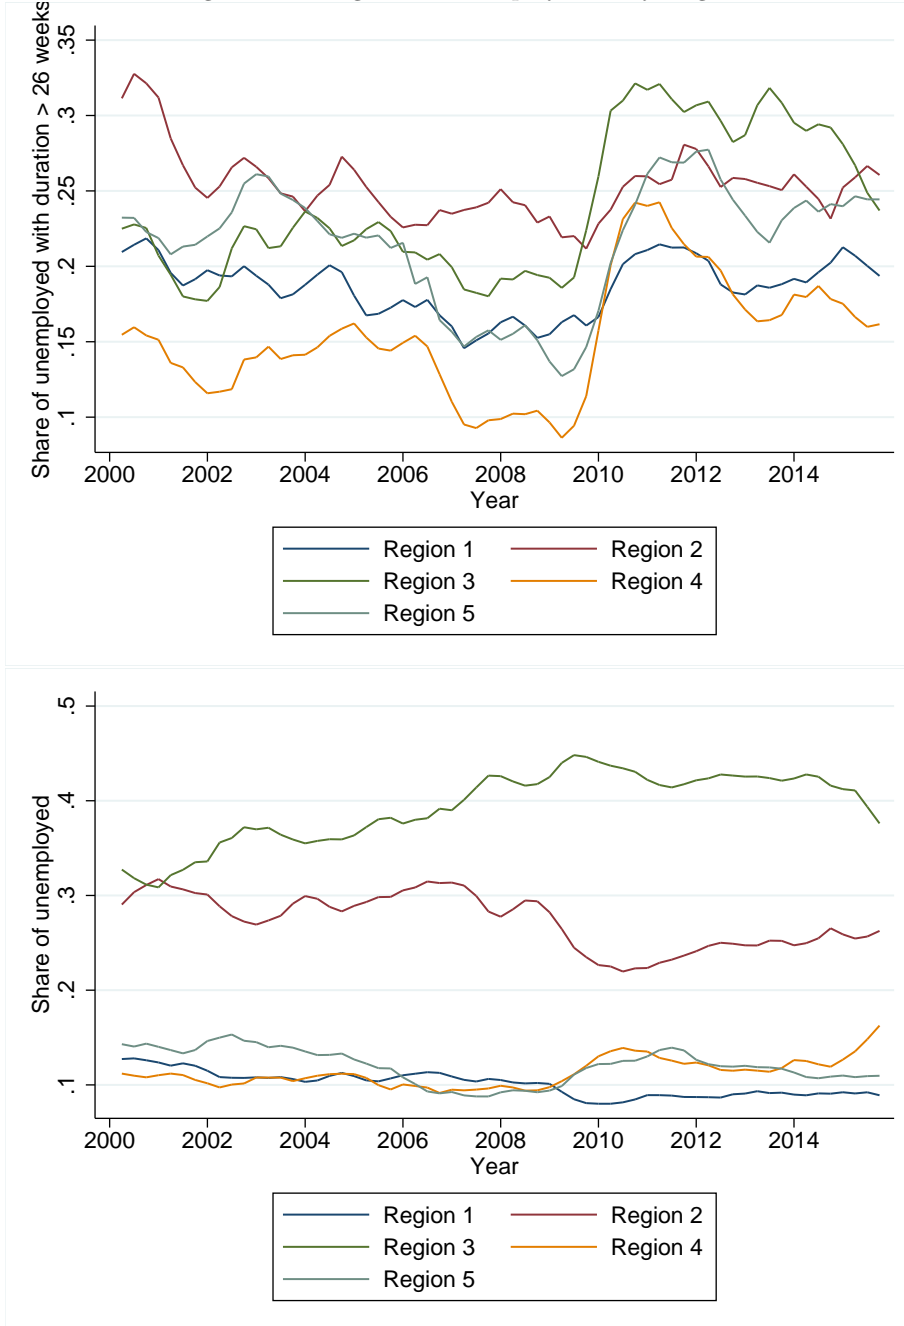


Figure 31: Long-Term Unemployment by Age



Figure 32: Long-Term Unemployment by Region



Note that the region codes are the following: Atlantic Provinces (1), Quebec (2), Ontario (3), Prairie Provinces (4), British Columbia (5).

Figure 33: Long-Term Joblessness for Non-Participants by Education

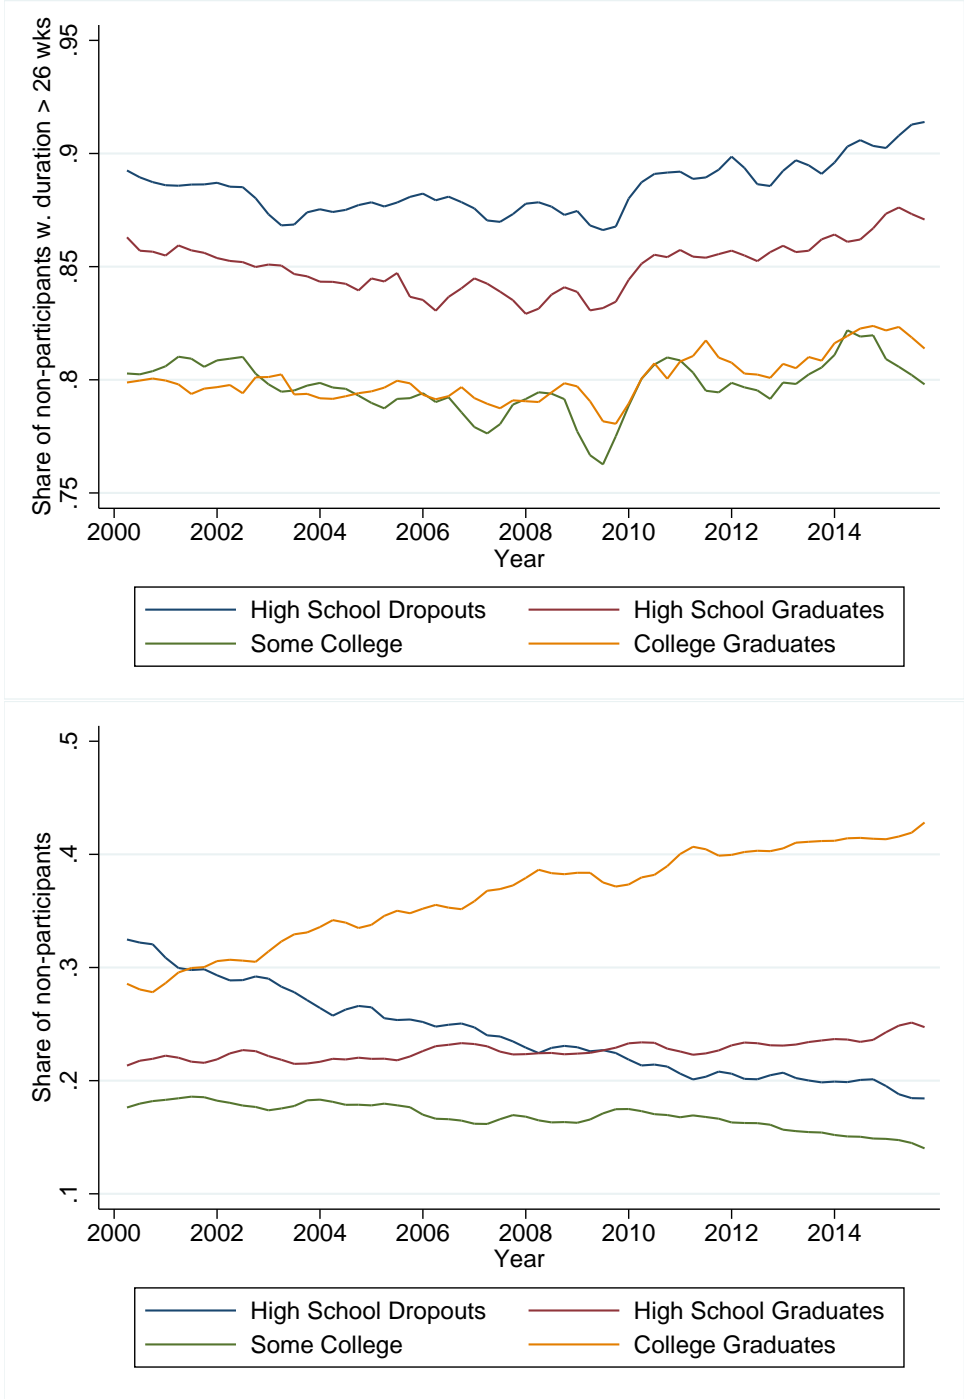


Figure 34: Long-Term Joblessness for Non-Participants by Gender

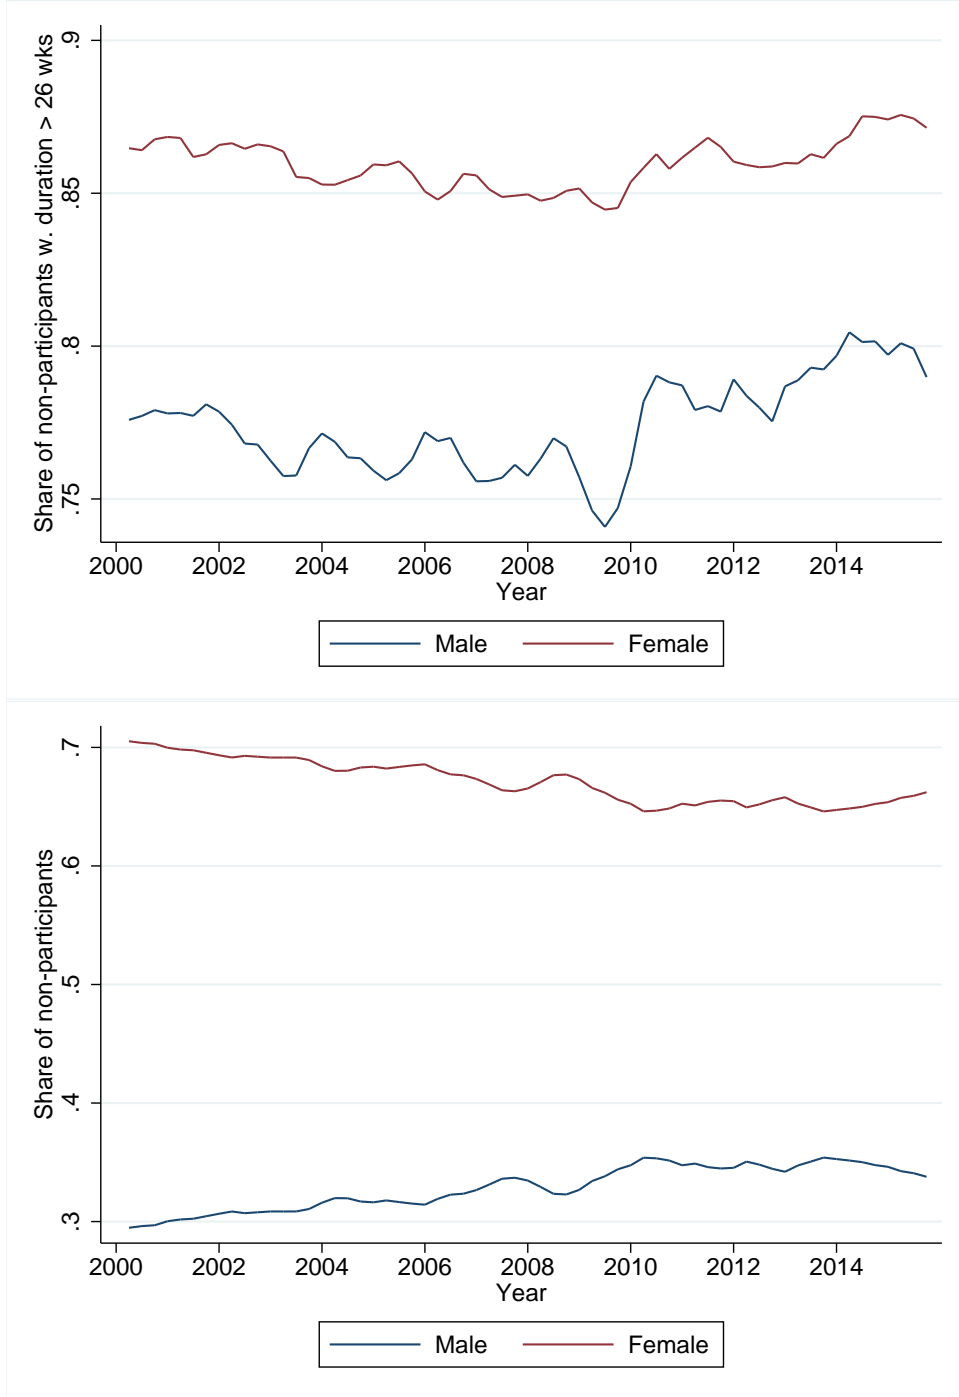
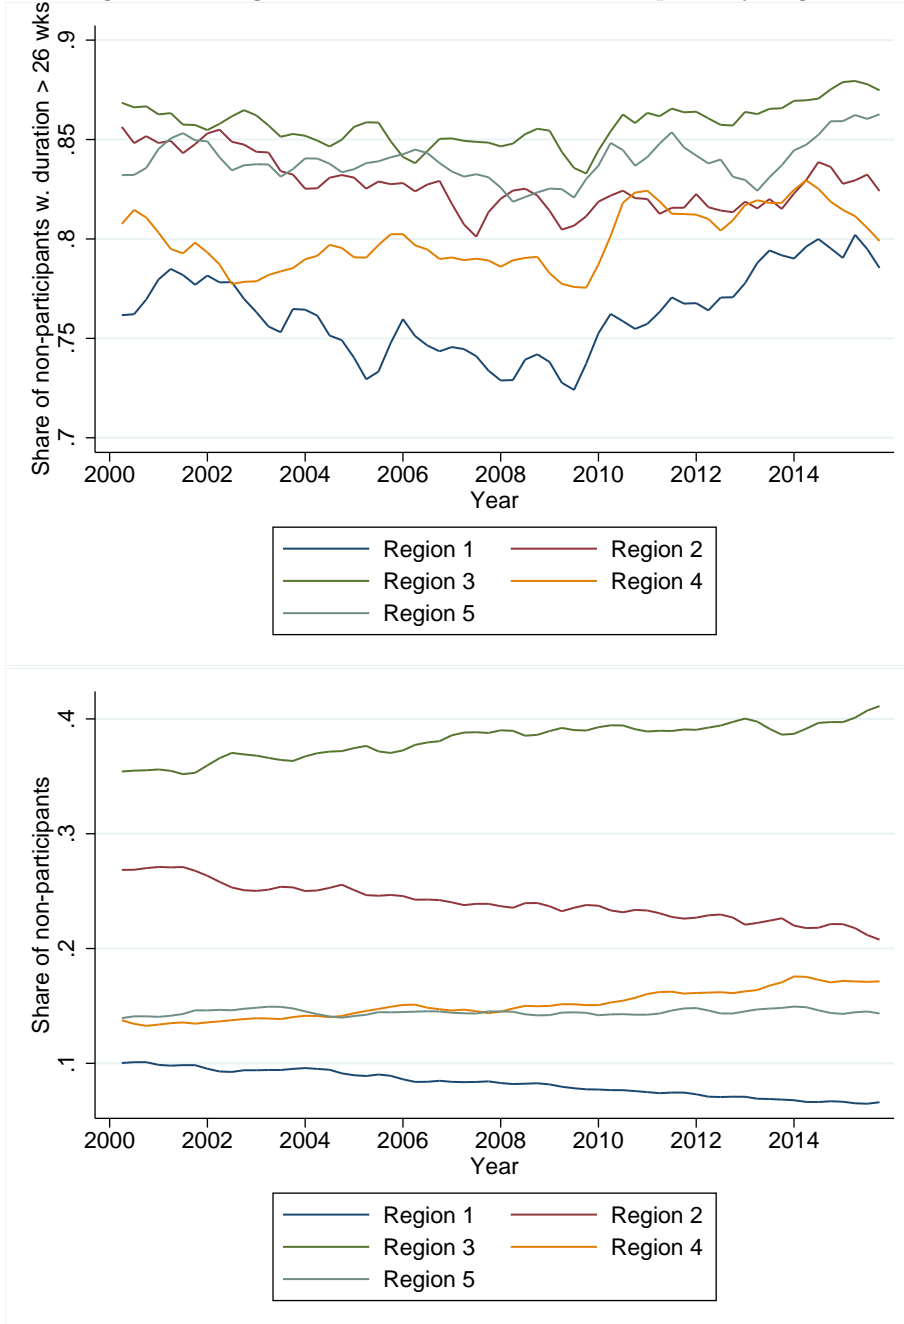


Figure 35: Long-Term Joblessness for Non-Participants by Age



Figure 36: Long-Term Joblessness for Non-Participants by Region



Note that the region codes are the following: Atlantic Provinces (1), Quebec (2), Ontario (3), Prairie Provinces (4), British Columbia (5).

Figure 37: job-finding Rate by Duration

