Canadian Labour Economics Forum

WORKING PAPER SERIES

Labour Force Transitions

Rui Castro (McGill University)
Fabian Lange (McGill University)
Markus Poschke (McGill University)

CLEF WP #78

Labor Force Transitions

Rui Castro, Fabian Lange, and Markus Poschke

August, 11th 2024

Abstract

Labor Force States and flows between are useful tools to model individual dynamics in the labor market. This chapter reviews recent literature uncovering substantial heterogeneity in transitions across Labor Force States. We review methods and results by replicating leading studies using Canadian data and relate our findings to important literatures on recall non-employment, duration dependence, and job ladders.

1 Introduction

Job loss and the ensuing risk of extended non-employment feature among the most significant economic risks workers are thought to face. One approach to measuring this risk is to define a set of labor force states meant to capture behaviors in the labor market, like work or job search, and to measure flow rates between them. The most common classification scheme uses survey responses, for example from the US

Current Population Survey (CPS), to divide the working age population into three groups, the employed (E), the unemployed (U), and those not-in-the-labor force (N).[1]

This three-part classification scheme only imperfectly captures the complexity of experiences in the labor market. Behavior within these states varies widely. For example, the unemployed might be more or less actively searching for a job, or they might wait to be recalled to a former job or for a new job to start. Similarly, both the employed and those not-in-the labor force might at times monitor the labor market for new opportunities. Moreover, individuals themselves permanently differ in ways not fully captured by point-in-time measures of labor force states. Understanding job loss risk, and labor market dynamics more generally, requires a framework featuring both an appropriate description of the labor market activities individuals undertake, as well as an appropriate description of their inherent differences.

The challenge of neatly classifying individuals into E, U, and N is illustrated by findings from the CPS re-interview program (Abowd & Zellner 1985, Poterba & Summers 1986). Within a week of the first interview, respondents were re-interviewed about activities during the same reference week. labor force states assigned by responses to both interviews often disagreed, in particular for U and N, suggesting that this distinction based on search activity is ambiguous. The observation by Elsby et al. (2015) that many respondents to the CPS report frequent switches between U and N within the same non-employment spell further adds to this impression.

However, the classic distinction between U and N is not without empirical content. Going back to Flinn & Heckman (1983), researchers have shown that job finding rates differ substantially across the two groups. Job finding rates also differ across

[1]Employed are those who performed at least one hour of paid work in a reference week. Unemployed are those who are not employed, available for work, and, with exceptions, actively searching for a job.

finer partitions of the non-employed. Jones & Riddell (1999, 2006) for instance report significantly higher job finding rates for the marginally attached as well as for discouraged workers.[2] In the same vein, Kudlyak & Lange (2014) and Hornstein et al. (2014) present evidence that respondents who switch between U and N within a non-employment spell are distinct from those who consistently report U or N throughout a given spell, suggesting heterogeneity within the U and N states.

Employment is not a homogeneous labor force state either. It is, for instance, well known that separation rates decline with job tenure (Mincer & Jovanovic 1981, Topel & Ward 1992). A more recent literature emphasizes heterogeneity between different types of employment ranked along a job ladder, where the different rungs capture vertically differentiated jobs. Workers prefer jobs on higher rungs and tend to move up the job ladder (Postel-Vinay & Robin 2002, Moscarini & Postel-Vinay 2018, Haltiwanger et al. 2018). Separation rates into non-employment or to other firms also vary systematically all along the ladder. Aggregating all employment in a single "employment" category thus obscures important heterogeneities in behavior.

A further challenge, noted already by Hall (1970) and Clark & Summers (1979), comes from permanent or at least persistent heterogeneity across individuals which extends beyond easily observable attributes like educational attainment. Ever since Lancaster (1979), a large literature has engaged with the difficulties permanent differences in job finding rates introduce when trying to estimate the causal effect of unemployment durations on job finding rates. One approach to identify genuine duration dependence (GDD) is to combine observable differences with functional form assumptions to separately identify GDD and dynamic selection. Nonparame-

[2] Marginally attached are those not-in-the-labor force who report that they desire a job, but are not currently searching. Discouraged workers are those who report that they are not searching because they believe no work suited for them is available in their area.

teric identification however requires data where individuals are observed for multiple spells (Heckman & Singer 1984*b*).

Thus, the traditional scheme of classifying the working age population into E, U, and N misses important variation in behaviors and across individuals. Over recent decades, evidence on these heterogeneities coming from survey panel data and from administrative sources has accumulated. This chapter engages with this literature. We focus particularly on a branch of the literature that exploits panel data with repeated spells in different labor force states to estimate transition models that allow for permanent individual-level heterogeneity and that allows for latent labor force states to capture variation in behavior otherwise obscured by observed labor force states.

Key to this literature is the increasing availability of panel data which allows measuring repeated employment and non-employment spells. One approach, exemplified by Shibata (2019), Hall & Kudlyak (2022) and Ahn et al. (2023), exploits the short panels of the CPS to estimate models with multiple types of individuals that transition across latent states following a Markov process. We refer to this approach as the 'short panel approach.'

A related approach, creatively termed the 'long panel approach', uses long panels derived either from surveys (Morchio 2020) or administrative data (Gregory et al. 2021). Long panels are well suited for identifying permanent individual heterogeneity, while it is hard to differentiate between transitory and permanent variation using short panels. Compared to short panels, long panels often suffer from a lower frequency (typically quarterly or even annual). Long panels derived from administrative data typically lack detailed information on demographics and activities like search. The difference in sample size and in the complexity of labor force histories between short and long panels furthermore means that the two approaches are also differentiated by the statistical methods employed.

This chapter begins with a review of this recent literature, which we summarize with the help of a simple unifying framework featuring type-specific Markov transitions across latent labor market states: the Hidden Markov Model. We aim to distil commonalities across these different approaches and their substantive findings, and to sketch their broader implications. To do so in a conherent manner, we focus on a key paper for each approach (Hall & Kudlyak (2022) for the short panel approach and Gregory et al. (2021) for the long panel approach) and replicate its core analysis using comparable Canadian data sources. We then synthesize broad lessons across the existing literature and our replications. We then turn to some more specific recent literatures on labor market flows in which latent states and heterogeneity loom large, namely those on duration dependence in labor market transitions, recall from unemployment, and job ladders, and use our estimates to illustrate the potential importance of unobservable heterogeneity for these topics.

We find that a broadly similar picture of individual dynamics in the labor market emerges from the two approaches and types of data sources. This same picture also describes dynamics in both the U.S. and Canada well. First, the labor market consists of highly differentiated types of individuals. For no type do flow rates across labor force states equal those obtained when averaging across the population. Failing to account for the heterogeneity in the population leads to misleading characterization of the dynamics and the risk individuals face in the labor market. Second, a vast majority of the population is almost always employed and faces little risk of job loss, especially of long duration. Such risks are instead concentrated in a minority of the population who experience many non-employment spells, spend significant time non-employed, account for most of aggregate non-employment, and earn somewhat lower wages when employed. These findings imply that the costs of uncertainty associated with labor market flows differ significantly across types.

Even though the two approaches agree in their main findings, they do differ in the role they ascribe to job finding and job separation rates in generating different labor market experiences across types in the population. Using the short panels from the Canadian labor Force Survey (C-LFS), we find, as do Hall & Kudlyak (2022) using the CPS, that types differ most strongly in job finding rates. Our analysis using long panel data by contrast suggests that types are differentiated mostly by job separation rates. We conjecture that these differences partially reflect the use of short vs long panels. Whether heterogeneity across types is mostly attributable to non-employment inflow or outflow rates remains an important open question.

We then describe connections to three important literatures on labor market dynamics, starting with the long literature on duration dependence in labor market flows, in particular job finding and job separation rates. We believe the long-panel data is better suited to identify type heterogeneity and thus turn to this data to revisit the still-contested question of the relative importance of GDD and dynamic selection in shaping the decline in job finding rates with unemployment duration. Because job finding rates differ little across types in the Canadian long panel data, we find a very limited role for dynamic selection. Rather, we ascribe most of the observed decline in job finding rates (after accounting for recall) to GDD within all types. Our finding of GDD in job finding rates contrasts with findings in the literature (Machin & Manning (1999) and others reviewed below) that dynamic selection based on type heterogeneity explains much of observed duration dependence in job finding.[3]

[3] We speculate that this difference arises because, when following the long-panel approach, we identify labor market types using differences in job finding and job separation rates, while the

About a third of non-employment spells, and in some countries more, end in recall to the original employer (Fujita & Moscarini 2017, Katz 1986, Kroft et al. 2019).[4] In addition, the duration structure of recall unemployment is very different from that of search unemployment. To analyze the dynamics of job finding thus necessarily requires engaging with recall non-employment. We find that recall accounts for a large share of job finding in the first year of the unemployment spell. Subsequently, the probability of recall becomes negligible, implying clear duration dependence in recall. As for job finding rates ending in new employment relationships, we find that most of the duration structure of recall reflects GDD, not differences across types. Those differences we find across types suggest that recall is relatively more prevalent among the more attached types - not only do these experience fewer non-employment spells, but their non-employment spells tend also to end more often in rejoining their prior employer.

Finally, we connect to the very active literature on job ladders. A job ladder generates wage growth, and also implies differences in labor market flows by job and employment tenure. It also helps explain persistent earnings losses after a layoff. Some recent work indicates that the ladder is particularly important for earnings growth of low-skill workers. Conversely, workers who experience more frequent job loss have greater difficulty climbing the ladder, and experience less earnings growth.

We examine whether differ labor market types, as identified using the long-panel approach, fare differently on the job ladder. Not surprisingly, the highly attached

literature on GDD exploits primarily job finding rates. That literature might thus amplify heterogeneities in job finding rates among the unemployed relative to our approach. It remains an open question how types estimated on specific transition rates compare to those estimated on all transitions jointly.

[4] Fujita & Moscarini and Kroft et al. note that individuals are often recalled not just out of unemployment but from non-participation N.

types tend to be found on the top rungs of the job ladder, while the less attached types are typically found at the bottom. This distribution of types along the ladder does not, however, stem primarily from differences in how types climb the ladder once they find employment. Rather, high types tend to join the ladder on higher rungs, and are then much more likely to remain there. Low types, by contrast, tend to join the ladder at the bottom, and then need to climb it to reach the higher rungs. At the same time, however, they face high separation rates at every rung. For these low types, the ladder might be better described as a greasy pole.

In Section 2, we begin by presenting a basic framework for modeling heterogeneity in labor force transitions. The recent literature is described in light of this framework. Section 3 contrasts findings in the literature pertaining to the US to our findings from Canadian data. Section 4 assesses the implications for how labor market risk as experienced by individuals is distributed in the population. The next three sections connect to the literatures on duration dependence in flows (Section 5), recall non-employment (Section 6), and job ladders (Section 7). We conclude in Section 8.

2 Heterogeneity in a Markovian World

2.1 A basic framework to model heterogeneity in labor force transitions

This section outlines a Hidden Markov Model, a basic structure that unifies the recent literature on heterogeneity in labor force transitions but maintains a Markovian structure. Agents transition across a set of distinct latent labor force states following a Markov process. Permanent heterogeneity in transitions across agents is captured by types that differ in transition probabilities across latent states. The latent states then map into a set of observed states. This mapping can but need not be one-to-one. It is for instance possible that multiple latent states map into a single observed state, and vice versa. With this structure, the transitions across observed states need not be Markovian anymore, allowing these models to fit basic observed patterns in the data.

Set-ups of the type discussed here prove to be very flexible in that the researcher has many degrees of freedom in specifying latent and observed states as well as the number of types. States (latent or observed) can include unemployment and inactivity of different durations (see Section 5), recall or temporary unemployment (Section 6), different rungs of a job ladder (Section 7), or other labor force states.

More formally, there are Ω types indexed by ω belonging to a set S_ω. The corresponding triplet for the latent states is denoted (K, k, S_k) and for the observed states (M, m, S_m). The matrix $\Gamma(\omega)$ maps latent to observed states. The probability distribution across types is given by a probability vector θ and is fixed over time.

Transitions between states are described by a Markov transition matrix $\Pi(\omega)$ with representative element $\Pi_{ij}(\omega)$ (row i, column j), which is the probability of moving from latent state i to latent state j for type ω. Refer to row k as $\Pi_k(\omega)$. These transition probabilities reflect transition determinants that can be both exogenous (e.g. involuntary job loss due to firm closure) and endogenous (e.g. job finding probability as a result of search effort, job acceptance probabilities, etc.). Different economic models will give rise to different restrictions on these transition matrices.[5][6]

[5] For surveys of the theoretical literature on job search, see Wright et al. (2021) and Rogerson et al. (2005).

[6] For identification of binary Markov models with unobserved heterogeneity consider Browning & Carro (2014). Shibata (2019) contains useful references regarding identification of Hidden Markov Models.

We consider stationary environments. The transition matrix induces an ergodic distribution across latent states.[7] Let $n_k(\omega)$ be the fraction of the population of type ω in state k, and let $\mathbf{n}(\omega)$ be the column vector $[n_1(\omega), n_2(\omega), ...n_K(\omega)]$. Naturally, $\sum_k n_k(\omega)$ equals one. This distribution follows the law of motion

$$\mathbf{n}'(\omega) = \Pi(\omega)^T \mathbf{n}(\omega)$$

where a prime refers to variables for "tomorrow", and superscript T denotes the transpose. In a stationary state, $\mathbf{n}'(\omega) = \mathbf{n}(\omega)$, which is the solution to the equation $(I - \Pi(\omega)^T)\mathbf{n}(\omega) = 0$, where I is the identity matrix. The steady state distribution $\mathbf{n}^*(\omega)$ is also the first eigenvector of the transition matrix $\Pi(\omega)^T$.

The matrix Γ mapping latent states k into observed states m then generates a type-specific ergodic distribution as well as a type-specific joint distribution across observed states.[8] Aggregating across types ω using the type-distribution vector θ delivers the corresponding population distributions.

Variations of the models of this type are differentiated by their restrictions on the transition matrices Π and measurement matrices Γ. It should be noted that many versions of these Hidden Markov Models are isomorphic. For instance, models with multiple types ω and type specific transition matrices $\Pi(\omega)$ are isomorphic to models with a single type and a block diagonal transition matrix. And, empirically, intermediate cases will be relevant where agents might be classified into types because

[7]For a set-up that explicitly considers time variation, see Ahn et al. (2023), who estimate a model of this type using 1980-2021 Current Population Survey data. They allow the transition matrices to vary by month, which induces time-varying distributions that can be used to study growth, business cycle, and seasonal dynamics. We abstract from this type of dynamics. Shibata (2019) also considers some examples with time variation.

[8]Most contributors to this literature specify this mapping to be deterministic, so that each latent state maps into one observed state. Ahn et al. (2023) allow for probabilistic mappings.

the transition matrix is close to but not strictly block diagonal.

Hidden Markov Models such as the one described above can be specialized in a variety of ways making use of different variables contained in panel surveys and imposing different restrictions to aid identifying the models as well as allowing for interpretations of the estimated structures. Models of this type are very flexible and can reproduce many empirical patterns that Markov models with homogeneous types and only observed states can't. Importantly, they can generate duration dependence in observed transition rates either by allowing for latent states such as short- and long-term employment (Hall & Kudlyak 2022) or short- and long-term unemployment (Ahn et al. 2023) and/or by allowing for type heterogeneity. We will discuss links to the literature on duration dependence in Section 5 below. These models can also accommodate job ladders, which we consider in Section 7 below.

2.2 Four Examples

Table 1 collects information on the number of types, observed states, and latent states in four recent papers that either closely (Ahn et al. (2023), Hall & Kudlyak (2022), and Shibata (2019)) or more loosely (Gregory et al. 2021) adhere to the structure outlined above.

The first three papers use the short panels from labor force surveys to estimate Hidden Markov Models of labor force transitions. Shibata (2019) considers a very flexible formulation with one type, 9 latent states, and the typical 3 observed states E, U, and N.

Similar to Shibata (2019), Hall & Kudlyak (2022) limit themselves to 3 observed states, but allow for five types and four latent states. The latent states they consider are U and N, which are both directly observed, as well as short-term and long-term employment. Short-term and long-term employment are not directly observed in the CPS and thus represent pure latent states.

By contrast, Ahn et al. (2023) exploit more fully the detailed information on labor force activities in the CPS. They allow for 29 observed states, including temporary layoffs, unemployment of different durations, and the information commonly used to classify respondents not in the labor force into marginally attached or discouraged workers. Contrary to Hall & Kudlyak (2022), they specify the latent states to contain a short- and long-term unemployment state, but do not distinguish between short- and long-term employment.

Gregory et al. (2021) (GMW hereafter), the final paper listed in Table 1, departs in its approach from the prior three papers. Empirically, it relies on a long panel from the US Census Longitudinal Employer-Household Dynamics program (LEHD), rather than the CPS. In this data, only two observed states can be distinguished: employment (E) and non-employment (NE). The authors do not estimate the underlying Markov model directly. They first estimate types and then, conditional on these types, estimate a fully-specified economic model. Rather than estimate their model, we will follow the alternative approach of exploring transitions conditional on types. From now on, when we refer to GMW, we refer to this alternative approach of first identifying the types in the data following Gregory et al. (2021), and then estimating type-specific transition matrices, without the need for latent states. This facilitates comparability with the aforementioned papers estimating Hidden Markov Models on short panels.

Table 1: Heterogeneity in Transitions: Approaches in Recent Literature

Paper Main Ingredients	Ahn et al. (2023)	Hall & Kudlyak (2022)	Shibata (2019)	Gregory et al. (2021)
Types	3 types ("primary", "secondary", and "tertiary")	5: 2 fixed types restricted to one state (all-E and all-N) and 3 mover types with flows across all states (high-E, high-U, high-N)	1 type	4: 3 estimated types + 1 sample restriction excluding individuals with nonemployment spells longer than 2 years
Latent States	4: E, short-term U, long-term U, N	4: U, N, short-term, and long-term E	9	N/A
Observed States	29 incl. part-time and full-time E, N by reason, U by duration, recall	3: E, U, N	3: E, U, N	2: E and NE

The table summarizes the states and types as specified in 4 recent papers describing heterogeneity in the labor market. E stands for employment, U stands for unemployment, N for non-participants, and NE for nonemployment (NE=U+N). Gregory et al. (2021) do not explicitly estimate a latent state model.

2.3 Estimation using Short and Long Panels

The papers listed in Table 1 differ in the type of data used for estimation. The first set of papers (Ahn et al. (2023), Hall & Kudlyak (2022), and Shibata (2019)) rely on short panels from the CPS. Similar data are collected by many statistical agencies around the world, notably the labor Force Survey (C-LFS) in Canada, which we analyze. Both the CPS and the C-LFS are described in more detail in Section 3. These labor force surveys provide high-frequency (typically monthly) observations and contain rich information about an individual's labor market activity, allowing in particular to distinguish between different types of non-employment, such as whether an individual is unemployed or not in the labor force. Comprehensive information about individual-level demographics is also usually available. The main drawback is that these panels tend to be quite short. The CPS for instance consists of two four-month panels, separated by an eight-month gap. As we discuss in Section 5 below, separating type heterogeneity from duration dependence generated by flows across latent spells intuitively requires data with multiple spells of the same individual in the same labor force state. In short panels, such repeated spells are rare, making it hard to separate type heterogeneity from state dependence.

A second group of papers relies on long panel sources from administrative data, often based on tax records, from which decade-long employment histories can be extracted. Examples are the US Longitudinal Employer-Household Dynamics (LEHD) and the Canadian Employer Employee Dynamics Database (CEEDD), in particular the extract contained in the Business Employee Analytical Microdata (BEAM), and are described in more detail in Section 3. These long panels are better suited to isolate individual heterogeneity, and to study issues relating to job ladders and life-cycle dynamics. Several of these data sources also link workers to firms, allowing for inference about job quality as a function of firm-level characteristics. The main drawback of these panels is coarse information, namely infrequent (quarterly or less) observations, which is bound to miss labor flows, and scarce data on demographic characteristics or labor market activities (the type of non-employment is typically not observed). Our main reference using this type of data is Gregory et al. (2021), who analyze data from the LEHD.

In short panels, the number of possible histories is relatively small. For example, the CPS with 3 observed states allows only for $3^8 = 6,561$ different histories. The dynamics in the CPS are fully described by the probability vector over these histories. Using the model-based distribution across these histories, Hall & Kudlyak (2022) construct the likelihood and estimate the parameters using a full maximum likelihood estimator. Shibata (2019) and Ahn et al. (2023) instead rely on the iterative Expectation-Maximization (EM) algorithm for estimation.[9]

With long panels, likelihood-based methods rapidly become computationally intensive. Even with just two employment outcomes (employment and non-employment), when T is at least an order of magnitude larger than in short panels, as is usually the case, then jointly estimating individual heterogeneity and type-specific model parameters governing individual employment dynamics becomes unfeasible. For example, just a six year panel of quarterly data would allow for 16.8 million histories. Bonhomme et al. (2022) provide an alternative two-stage grouped-fixed effect (GFE) procedure that is particularly useful in this regard, as it decouples the two estimation steps. The first stage uncovers and discretizes latent, possibly continuous individual heterogeneity into a small number of discrete types. This is done using the *k-means*

[9] In these papers, pinning down the number of unobserved types/states is based upon a combination of restrictions imposed by model identification, such as the one provided by Allman et al. (2009), and on sensitivity exercises with respect to the number of types/states, associated to a likelihood-based measure of model fit.

clustering algorithm.[10] The algorithm identifies an individual as being of a given type if her employment dynamics induces moments that are mostly similar to those of other individuals already identified to be of that same type. The second stage then separately estimates the model of interest, conditional on the group-specific heterogeneity identified in the first stage.

Gregory et al. (2021) follow the latter approach. To estimate types, they generate for each agent a set of moments summarizing their careers, based on 18 years of individual data. These moments include, for example, the fraction of unemployment spells less than 3 months long. They then use *k-means* clustering on these moments to uncover underlying types that capture the heterogeneity in employment histories.[11,12]

2.4 Basic Findings

Table 2 lists the substantive findings from the papers reviewed in Table 1. All four papers agree that some combination of latent states and type heterogeneity is crucial for understanding aggregate labor market transitions. The three papers emphasizing type heterogeneity, (Ahn et al. (2023), Hall & Kudlyak (2022), and Gregory et al. (2021)) also share the very basic finding that individual transitions are very different

[10] See Bonhomme & Manresa (2015) for a description of methods available to implement the clustering.

[11] After separating out agents with non-employment spells lasting more than 2 years, they allocate the remaining individuals to three estimated types, making it four types overall. The number of estimated types is chosen using the cross-validation approach proposed by Wang (2010). Note also that the clustering method attributes a type to each individual in the sample. Likelihood based methods instead directly estimate the HMM parameters, and only identify an individual's type in a probabilistic sense.

[12] Morchio (2020) is a related paper using a long panel, the NLSY79 survey, to estimate a model describing how individual employment risk varies over the life-cycle. In contrast to Gregory et al. (2021), this paper restricts individual unobserved heterogeneity a priori to two discrete types. However, like the second stage of Gregory et al.'s (2021) approach, it estimates a fully-specified model conditional on those types.

across types in the population, and that none of the types resembles the aggregate behavior. Most people spend most of their time employed, and experience very few transitions and very few non-employment spells, which also tend to be short-lived. They all report around 60% of the respective populations studied to be of such a strongly attached, stable type. Most of the observed flows and a substantial proportion of non-employment are accounted for by an unstable type that only makes up a small fraction of the population.

All four papers investigate the relationship between demographics and heterogeneity in transitions as summarized by the latent types that describe the cross-sectional heterogeneity in transitions. Table 2 summarizes their findings, which are broadly consistent across the four studies. Race, gender, education, and age typically strongly predict type membership (latent labor state attachment in the case of Shibata (2019)). However, these demographic predictors tend to account for only a relatively small fraction of the cross-sectional variation in type membership as measured by statistical fit criteria such as Pseudo-R^2 measures. In addition, the variation in labor force states across types uncovered by these methods is typically far larger than that across demographic characteristics. For example, Table 3 shows that in the CPS data used by Hall & Kudlyak (2022), the difference in the employment rate between men and women is 14.2 percentage points. By contrast, the difference in employment rates between the High-E and High-N types is 53 percentage points, with similarly stark differences across the other types.

In the following, we focus on unobserved heterogeneity and labor market dynamics. There is a vast literature on demographic differences in labor market outcomes, and we are not in the position to do this literature justice. As a starting point, we refer the reader to the excellent survey by Altonji & Blank (1999) in an earlier

Handbook of this series.

Table 2: Heterogeneity in Transitions: Substantive Findings in Recent Literature

Findings / Paper	Ahn et al. (2023)	Hall & Kudlyak (2022)	Shibata (2019)	Gregory et al. (2021)
Heterogeneity	55% in primary labor market with very few and short NE spells. Most flows accounted for by 14% in secondary market with frequent, long U spells. Tertiary sector mostly in N.	Two-thirds of population experience very few flows. 59% of aggregate U accounted for by 5% of population.	Relative to baseline FOM model with observable states, Hidden Markov Model predicts aggregate transitions significantly better.	57% of population strongly attached: almost always in E, with most E spells longer than 2 years, and NE spells shorter than 1 quarter. 17% population weakly attached: one-third of the time in NE, with most NE spells longer than 1 quarter and very few E spells longer than 2 years.
Demographics	Education, race, gender, and age strongly predict labor force types. However, these variables capture only a relatively small fraction of the cross-sectional variation in type membership.	By gender, type specific transition rates are similar but distribution across types differs significantly.	Men, married individuals, and the highly educated more likely to be highly attached. Nevertheless, significant heterogeneity across latent states exists within demographic groups.	Race, gender, and education are very poor predictors of type heterogeneity. However, white, males, and the highly educated are more likely to be of the strongly attached type.
Data	US CPS 1980-2021, 15 and older.	US CPS 2014-2017, 25-54 years old.	US CPS 1976 - 7/2014, 16-65 years old.	US LEHD 1997 - 2014

Note: The table summarizes the basic findings in 4 recent papers describing heterogeneity in the labor market. E stands for employment, U stands for unemployment, N for non-participants, and NE for nonemployment (NE=U+N).

3 Canada and the US – Short and Long Panels

We now contrast findings from the US-centered literature with own calculations based on Canadian data. Our approach is to replicate one short-panel and one long-panel study from the US using Canadian data. In this way, we aim to understand which results are robust across methodologies and which results might generalize across countries. As examples from the US literature, we use Hall & Kudlyak (2022) for the short panel literature and Gregory et al. (2021) for the long panel. Up to constraints imposed by data, we replicate their approaches on Canadian short panel data from the C-LFS and long panel administrative data from the Business Employee Analytical Microdata (BEAM).

3.1 Data

3.1.1 Short panel data

Hall & Kudlyak (2022) estimate their model using data from 25-54 year old respondents to the CPS. The observed states in Hall & Kudlyak (2022), are E, U, N. CPS respondents are surveyed monthly for four months and, after a break of eight months, for another four months.

To estimate their model, HK maximize the likelihood of the observed eight-month histories accounting for the gap in the middle. With more states as in Ahn et al. (2023), the number of parameters expands rapidly, especially given that the mapping of latent to observed states needs to be estimated and is not imposed on a priori grounds. This increases the computational burden and leads them to use the expectation-maximization algorithm. We refer the interested reader to the original papers for details on implementing the estimators.

We estimate HK's model on 25 to 54 year olds from the C-LFS. The C-LFS follows the (somewhat loose) guidelines of the International Labor Organization for defining the three labor force states E, U, and N. This allows for some, but not 100% comparability in measures across countries that follow these guidelines.[13]

The C-LFS also differs from the CPS in the sequencing of interviews. Instead of following the CPS's 4-8-4 pattern, it surveys respondents for 6 consecutive months. Thus, we have a total of $729 = 3^6$ possible employment histories, compared to the $6,561 = 3^8$ in the CPS.

Our analysis period covers 2001 to 2019 to align with the period used for the Canadian long panel described next. HK analyze data from 2014-2017, a period of relative calm, whereas our study period includes the Great Recession. Even though labor markets in Canada exhibited much less turmoil during the Great Recession than the US (Kroft et al. 2019), some of the differences uncovered across the US and Canada may arise from the different the time-periods studied.

3.1.2 Long panel data

In recent decades, researchers started using administrative panels covering close to the entire population and collected across long periods, as required for the long

[13] Canadian measurement conventions result in unemployment rates about 1 percentage point higher than US measurement conventions (Bender 2016). The higher rate is mostly due to the fact that Canada computes aggregate labor force statistics on the population aged 15 and over (compared to 16 and over in the US) and includes "future starts" (non-employed with jobs lined up to start within 4 weeks) among the unemployed, unlike the US. The former difference does not affect our analysis. To achieve greater comparability, we exclude future starts from U and treat them as N.

panel approach. In the US, the best known are from the US Census's Longitudinal Employer-Household Dynamics (LEHD) program and are based on records collected by the unemployment insurance programs run by the various states. GMW use the LEHD for 1997 to 2014 for a 2% random sample of individuals from 17 states, including California, Illinois and Texas. The LEHD does not include information on work that is not covered by unemployment insurance, like employment in the military or the US federal government, or self-employment, and is reported at quarterly frequency. It does also not contain information on time worked or on the nature of non-employment episodes (U or N). Furthermore, the data records earnings employees receive from a given employer in a quarter. Non-employment is not directly observed and needs to be imputed. This situation is typical (though not universal) for administrative data sets and can be empirically consequential, as it necessitates ad hoc judgements that are to some degree arbitrary. For any quarter without earnings, GMW impute a full quarter of NE; they impute NE of half a quarter when an individual has earnings from only one employer in one quarter and only from a different one in the following quarter, or when an individual has earnings from two employers in a quarter, but their total is less than the earnings in both surrounding quarters.

To cluster the population, GMW compute the following moments based on individual employment histories: (i) the fraction of time spent non-employed, (ii) the number of different jobs relative to the number of quarters in the data, (iii) the duration distribution of employment spells (share <1 quarter, 1 to 4 quarters, 4 to 8 quarters, and more than 8 quarters), and (iv) the duration distribution of non-employment spells (share <1 quarter, 1 to 4 quarters, or 4 to 8 quarters). They exclude individuals with any NE spell longer than 8 quarters. To construct these moments they need to know the spell duration up to 2 years for any spell ongoing between January 1999 and December 2012. GMW thus consider the two years before 1999 and after 2012 as "reference periods" to compute spell durations of spells ongoing at either end of the study period.

The BEAM is the corresponding Canadian data based on the Canadian Employer-Employee Dynamics Database (CEEDD). These anonymized data derive from individual and firm tax filings. The individual tax files provide data on earnings, as well as "Records of Employment" (ROEs) that employers are required to file whenever an employment relationship ends. The firm data includes information to identify mergers and splits of companies, which we also use to determine the length of employment spells. We study a random 1% sample of paid employees aged 25 to 54 who were issued a tax slip by an employer detailing earnings in any year between 2001 and 2019.[14]

Based on the ROEs, which record the start and end dates of employment relationships, we can, in principle, determine the history of employment spells. In practice, some ROEs are missing. This is by construction for all spells still ongoing at the end of the study period. However, sometimes an ROE that should exist is simply missing. In both cases, we impute spells based on the available information.[15]

The result is a data set with jobs held in each calendar month subject to imputation error, just as in the LEHD data. We assign main jobs using earnings and

[14] In Canada these tax slips are known as T4s, equivalent to the US W-2 form.

[15] In the first case, we impute the spell start month (we observe the year) by comparing earnings in the first year of the spell to full-year earnings in the next year of the spell. In the second case, we impute both the beginning and end month of the spell in this way. When a spell does not cover a full calendar year, we use earnings in neighboring spells. Private conversation with Statistics Canada employees also suggests that end-dates of employment spells derived from ROEs are reliable but that start-dates are measured with more error.

impute non-employment in months not covered by any employment spell.[16]

We follow GMW in that we use the first and last two years of the study-period (2001-2019) to determine the length of spells (up to two years) for all spells observed between 2003-2017. We then construct the same moments as do GMW using 2003-2017 and k-means cluster the resulting data-set of moments. This approach assigns a type to each individual in the sample. We deviate from GMW in that they exclude all individuals with non-employment spells exceeding 2 years, given their focus on business cycle fluctuations. As we aim to describe the entirety of the population, we do not exclude those with long non-employment spells from the sample.

3.2 Descriptions of Labor Markets

3.2.1 Prevalence and basic characteristics of labor market types

Table 3 shows the population proportions of five labor market types estimated on short panel data. The estimates by HK using the CPS are on the left, and our estimates using the C-LFS are on the right. The top panel shows information for women, and the bottom panel for men. This table also shows the fraction of time each type spends in each observable state, as implied by the ergodic distribution of the stochastic process.

Table 4 shows corresponding estimates using long panel data, with GMW on the left and our estimates using the Canadian BEAM on the right. The top panel shows results for women and men combined, following GMW's reporting. The bottom panel shows separate results by gender from the BEAM. Given the type estimates,

[16] We treat spells in the same firm that are interrupted for two months or less as a continuous employment spell, but use the non-employment spell when calculating distributional statistics related to non-employment. We also treat spells in firms linked by mergers or acquisitions as continuous.

the table also shows the fraction of time that each type spends in each observable state.

Type	US (CPS, Hall-Kudlyak)				Canada (LFS)			
	Pop.	Fraction time in			Pop.	Fraction time in		
	share	E	U	N	share	E	U	N
Women:								
All-E	50.2	100	0	0	52.0	100	0	0
High-E	15.0	89	4	7	18.5	90	4	7
High-U	6.0	57	28	15	9.8	64	27	9
High-N	15.4	34	4	63	11.8	32	5	64
All-N	13.4	0	0	100	7.8	0	0	100
Aggregate (data)		72.1	2.8	25.1		78.0	3.9	17.4
Men:								
All-E	66.4	100	0	0	62.4	100	0	0
High-E	15.2	90	5	5	16.1	89	7	4
High-U	5.4	57	34	10	9.9	63	32	5
High-N	7.2	43	7	51	8.3	51	6	42
All-N	5.7	0	0	100	3.3	0	0	100
Aggregate (data)		86.3	3.1	10.6		86.5	4.7	8.0

Table 3: Labor market states by type, short panels

Short panels. In their estimation, HK pre-specify that there are two types who are either always employed or always out of the labor force, and three "mover" types consisting of individuals who move across states. We take the same approach in the C-LFS data.

Table 3 shows that the most numerous type comprises half of all women, and almost two thirds of men. This is the always employed pre-specfied type, which HK refer to as "All-E". Similarly, the group pre-specified to be always out of the labor force, "All-N", accounts for about thirteen percent of women and five percent of men.[17]

In addition to these two types, who are observed in only one state, there are three types who move across states, but are predominantly observed in one of them. The second largest group, which accounts for about one sixth of women and men, consists of individuals who undergo a short or no non-employment spell, and correspondingly are employed for the entire or almost the entire time they are observed. HK refer to this type as "High-E".

Finally, there are two types of individuals who spend large fractions of time in either unemployment or non-employment. "High-U" individuals account for five to six percent of the sample in the US, and almost ten percent in Canada, which experienced more unemployment during the sample period. In contrast, while 15% of US women and 7% of men are classifed as "High-N", these shares are around 8 and 3%, respectively, in Canada, reflecting the higher participation rate in Canada.

"High-U" individuals spend about 30% of their time in unemployment, work around 60 percent of the time, and spend little time out of the labor force. "High-N" men work only slightly less: about half the time. This difference is larger for High-N women, who only work about half as much as High-U women. Both differ from their High-U counterparts in that their non-working time is spent almost entirely out of the labor force, and not in unemployment.

[17] The shares of these two groups in the population are estimated. The observed proportion of survey respondents who are always employed (non-employed) includes all All-E (All-NE) types as well as some individuals from other types who happen to experience long employment (non-employment) spells. Consequently, some High-E individuals may be observed as employed throughout their entire survey period in the CPS.

Overall, around 65 (US) to 70 (Canada) percent of women and about eighty percent of men belong to either the All-E or the High-E type, and are almost always observed in employment. The High-U, High-N and All-N types account for the remaining individuals. They spend large shares of time in unemployment or out of the labor force.

Comparing rows, it is clear that none of the types spends anywhere close to population average time in non-employment (about 13% for men and 21-28% for women, see bottom row). Most people spend very little time in unemployment or out of the labor force since the two groups accounting for the majority of the population, All-E and High-E, spend no to very little time in these states. But High-U, High-N and All-N individuals actually spend a large fraction of their time non-employed. Because they make up a minority of the population, the aggregate time spent in either unemployment or out of the labor force is nonetheless small, dominated by the first two groups.

As already transpired from this discussion, these results are similar by gender, although the model is estimated separately for women and men. For both genders, there are three mover types, and types have roughly similar population proportions and characteristics. The main difference between genders is that a greater proportion of women are High-N or All-N. As a result, the share of All-E women is smaller than that of men. The shares of High-E and High-U women are very similar to those for men. These differences reflect the higher participation rates of men in both countries.

Results for the US and Canada are strikingly similar, both in terms of type characteristics and type proportions. The two main difference is a smaller share of All-N and a larger share High-U in Canada, reflecting greater participation but also slightly higher unemployment.

What feature of the data and/or model generates the large share of individuals with very little mobility? The non-mover types help the model to generate not only the large fractions of individuals who are employed for the entire time in the sample, but also the decreasing separation hazards from employment, which is the focus of Section 5. The model does allow for short- and long-term employment but does not allow for separate short- and long-term unemployment states.[18]

At the same time, the large share of All-E types stems from the large fraction of respondents who happen to be always employed in a short panel such as the CPS or the C-LFS. We think this to be misleading. As we show below, the fraction of individuals who are always employed is smaller in the long panels. We conjecture that the estimated fraction of All-E relative to High-E types would be smaller if we estimated the HK model on a longer panel. As a result, the distinction between the All-E and High-E types is perhaps to an extent a function of survey characteristics. For this reason, it is useful to consider these two types together rather than as separate types, at least when referring to short panels.

Long panels. Table 4 displays results on labor force dynamics from the long panels. The top panel presents statistics for both genders combined, comparing the GMW estimation using U.S. data with our own analysis based on the Canadian BEAM dataset. The bottom panel then breaks the results from Canada out by men and women.

GMW exclude individuals with a non-employment spell longer than two years to focus on a highly attached sample, and estimate three types which they call the αs,

[18] Ahn (2023) allow for short- and long-term unemployment but do not allow for short- and long-term employment. They as well generate substantial heterogeneity in types with large fractions of individuals nearly always employed.

Panel A: Women and men:								
	US (LEHD, GMW)				Canada (BEAM)			
Type	Pop. share		Fraction time in		Share in Pop.		Fraction time in	
	overall	excl. δ	E	NE	overall	excl. δ	E	NE
α		57	96.4	3.6	40	51	96.9	3.1
β		26	90.4	9.6	14	18	91.6	8.4
γ		17	70.8	29.2	24	31	81.6	18.4
δ	excluded by construction				22		55.2	44.8
Total					100		83.3	16.7
Panel B: By gender:								
	Canada, women (BEAM)				Canada, men (BEAM)			
Type	Pop. share		Fraction time in		Share in Pop.		Fraction time in	
	overall	excl. δ	E	NE	overall	excl. δ	E	NE
α	38	48	95.8	4.2	41	52	97.8	2.2
β	15	19	90.2	9.8	14	17	92.9	7.1
γ	22	28	81.7	18.3	25	31	81.4	18.6
δ	24		54.9	45.1	20		55.5	44.5
Total	100		81.8	18.2	100		84.6	15.4

Table 4: Labor market states by type, long panels

the βs, and the γs.[19]

GMW find that the αs make up more than half of the sample. A second group, the βs, accounts for a fifth to a quarter of the sample, and is also employed more than ninety percent of the time. Our findings from Canada are similar, with a majority

[19] Unemployment duration moments cannot be computed for individuals who are always employed. GMW thus exclude them when they use k-means clustering to estimate types and subsequently include them in the type with the highest employment propensity, the αs. The statistics in Table 4 for the αs therefore include both those estimated to be of this type, and those in the sample who are always employed. We proceed in the same way with the Canadian data.

belonging to types who are almost always employed, even though this majority is smaller than in the US.

Instead, in Canada, the third group, the γs, is substantially larger. This group is significantly more often non-employed. In the US, this group is small, at about a sixth of the sample, and its members are non-employed almost thirty percent of the time. In Canadian data, they make up thirty percent of the sample that excludes δs and is thus directly comparable with GMW. This difference in shares exaggerates the differences between the US and Canada however since the γ's are more similar to the other types in Canada. Nonetheless, it is very distinct, as its non-employment rate almost reaches twenty percent.[20]

For Canada, Table 4 also reports statistics for the group of less attached individuals, here labelled δ, who experience one or more non-employment spells of more than two years. GMW exclude these individuals from their analysis, and no moments are reported. This group is sizable in Canada, accounting for more than a fifth of the population. This group splits its time about equally between employment and non-employment. Since this is a large group of individuals with a non-negligible contribution to aggregate employment, we find it pertinent to include them.[21]

We present results by gender in the bottom panel of Table 4.[22] This reveals very small gender differences. The largest ones are the higher non-employment rate of female αs and βs, and a slightly larger population share of δs. These factors jointly account for the slightly lower employment rate of women in the population.

[20]One interpretation is that the k-means algorithm partitions a continuous type space differently in Canada.

[21]For the Canadian data we report type shares in the whole population and, for comparability with GMW, shares in the sample that exclude δs.

[22]Types are estimated on the pooled sample of women and men.

Taking stock, the long panel paints a picture of labor market heterogeneity which is broadly consistent with the short panel. Like in the short panel, more than three quarters of the population belong to types with employment and non-employment rates very different from the population average. The average hides significant variation.

Long panels do, however, point to a much larger share of the population in the low attachment groups. In Canada, we find that the low attachment groups ($\gamma + \delta$) account for about 45% of the population when using the long panel. By contrast, the three low attachment groups that experience significant non-employment in the short panel (High-U, High-N, and All-N) together account for only about 22% of men and 29% of women. Yet, those identified to be less attached using the short panel are significantly less well attached to the labor market.

Interestingly, the share of individuals classified as δs by the long panel method corresponds roughly to the three low attachment groups obtained using HK on the C-LFS, and their employment shares are relatively close to those found by the short panel method. This again highlights the relevance of including this group in the analysis.

One should keep in mind, however, that comparisons across the short and long panels are not trivial, given that the groups identified by the different approaches do not closely map into each other. The approaches differ not only due to different samples, but also different estimation methods.

3.2.2 Contributions to aggregate employment and non-employment

Since types differ strongly in their employment rates, they contribute in very different ways to aggregate employment and non-employment. Table 5 shows the contribution of each type to aggregate employment and non-employment, based on the type's population proportion and mean time spent in E, U and N (or NE for the long panels). The top panel shows this for the short panels. The most striking finding here is that despite their small population share of only around twenty percent for men and thirty percent for women, the low attachment types (High-U, High-N and All-N) account for a disproportionate share of total non-employment. High-U individuals account for 60 to 70 percent of aggregate unemployment, while High-N and All-N individuals combined account for 85 to 90 percent of those out of the labor force. That is, almost all non-employment in the population is attributable to three groups that make up only a minority of the population.

The pattern is similar though somewhat less extreme in the Canadian BEAM long panel data set. Here, γs and δs, although a minority of the population, account for 83 (women) to 87 (men) percent of non-employment.

3.2.3 Flows between employment and non-employment by type

Table 5 also shows estimated flow rates by type, aggregated to flows between employment and non-employment for compactness.

Short panels. By definition, flow rates for non-movers are zero and we can only estimate flow rates for non-mover types, All-E and All-N. These estimates reveal that mover types differ in both job finding and separation rates.

The most glaring finding obtained using the HK model on short panels is that job finding rates vary substantially between types, whereas job loss rates are much more similar. With more than fifty percent per month, job finding rates (from NE to E) for High-E types are almost an order of magnitude larger than the rates of five to twelve percent experienced by High-U or High-N types.

Flows from E to NE differ much less across mover types. They are generally largest for High-N types, with a particularly large gap to the other types among women. We believe that the small differences in job separation rates across mover types partially results from the short sample length and from dividing highly attached individuals into High-E and All-E types. The E-NE rate for High-E and All-E individuals combined is only 1.2 percent (1.6 percent) for men (women) in US data, and 1.4 (1.8) percent for men (women) in Canadian data.[23]

In summary, we can say that types as identified by the short panel method differ significantly in their flow rates across states, particularly for flow rates from NE to E. Differences across types are so substantial that average flow rates, as obtained when estimating the HK model for a single type (ignoring heterogeneity, all else the same), are not well represented by any specific type.

Long panels. Maybe the largest differences across the short and long panel methods are evident when examining flow rates (these are measured directly in long panels). Whereas in the short panels, we found employment outflow rates to be quite similar across types, this is not true for the long panel. Among men, employment outflow rates for γs and δs are an order of magnitude larger than those for αs, with βs falling in between. Among women, the employment outflow rate of γs and δs are five times as large as those of αs. In contrast, employment inflow rates

[23]In a short sample, it is almost unavoidable for the High-E type to have a high separation rate. For example, if after five periods of E, a High-E individual transitions to NE at the end of the six-month spell in the C-LFS, this immediately translates into an E-NE rate of one fifth for that individual. (This high rate is tempered by the presence of other High-E individuals who spend all periods employed.)

are significantly lower among δs – not surprising given that this type is defined by having a long NE spell – but similar for the other three types. Thus, the long panel method identifies larger differences in employment outflow rates as opposed to job finding rates, whereas the short panel methods find the reverse pattern. It is unclear what gives rise to these differences across methods and data. As we just mentioned, we suspect that the short panel method tends to concentrate flows among highly attached individuals in the High-E and that the High-E and All-E group are not quite as distinct as the short panel method suggests. This results in large flow rates in either direction (E to NE and vice versa) in the High-E type, accentuating differences in flow rates between this and the other mover types for NE to E flows, while attenuating those in E to NE flows.

In long panels, just like in the short panel, type-specific flow rates do differ strongly from aggregate flow rates for all types. The high employment outflow rates of γs and δs stand out, and essentially drive their much higher non-employment rates compared to the αs and βs. More specifically, from the cross-state flow accounting identity, the ratio of the average time spent in E vs NE from Table 4 equals, at least approximately, the ratio of the job finding rate to the separation rate from Table 5. For a case in point, compare γ and α men in Canada. This ratio is only 4.6 for the γs, whereas it is much higher for the αs, 43.3. Given that the job finding rates are relatively similar between the two types, we conclude that cross-type differences in separation rates are indeed the main driver behind the large differences in non-employment rates.

3.2.4 State duration by labor market type

We next turn to differences in the duration of different employment states by type. Table 6 shows the duration distribution by type, for both short and long panel methods. This table compares results reported for the US data with our results based on the Canadian data broken down by gender for the short panel and aggregate across gender for the long panel method since GMW do not report gender specific statistics. Table 7 provides results for the long panel method by gender, using the Canadian data.

Short panels. Panel A of Table 6 shows the mean duration of employment and non-employment spells in the CPS and the C-LFS implied by the estimated type-specific transition matrices. By construction, flow rates out of their state are zero for All-E and All-N individuals and durations in the other state can not be calculated.

Mover types display variation in the mean duration of both E and NE spells. Throughout, High-N types have shorter employment spells and longer non-employment spells. High-U types differ from High-E ones not so much in their length of E spells, but much more in having significantly longer NE spells.

Results are very similar in US and Canadian data. In Canadian data, we can contrast these mean durations to those obtained when estimating the HK model for a single type. This reveals that all mover types have shorter E durations than the single type. This reflects the large weight of All-E types in the population. High-E types have much shorter NE durations, and High-U and High-N types have longer NE durations. That is, the population average describes no type, which echoes our analogous conclusion regarding heterogeneity in flow rates.

Long panels. Panel B of Table 6 shows the duration distribution of employment and non-employment spells in the LEHD as reported by GMW, and in the BEAM. This is for both genders combined. In the latter, we can also compute the mean duration of E and NE spells, as well as the duration distribution of spells for δs and across the whole population.

In the LEHD data, it is clear that αs and βs have longer E spells than γs. αs also have clearly shorter NE spells. Recall that, by definition, no group includes individuals with an NE spell longer than two years.

In the Canadian BEAM data, we see similar patterns. This is most salient in the column for mean spell duration. αs have much longer E spells than βs, who in turn have much longer E spells than γs or δs. The distribution of spell lengths reveals more detail. Almost all employment spells of αs last more than two years. For βs, this is still a very large share. For γs, in contrast, more than half of all E spells last a year or less. δs are roughly similar.

Regarding NE, and in contrast to the US, γs have the shortest spells on average in Canada. By construction, this is different for δs, since these all experience at least one NE spell of more than two years. The bottom line is that, here too, no type is well described by the average transition rates.

Table 7 shows that these patterns do not differ much by gender. The main difference is that women experience generally longer spells. This occurs for both E and NE spells, and for all types. This is a consequence both of fewer short spells and more long spells.

3.2.5 Relative wages

We conclude the description of the different labor market types by comparing their wages (short panel) or earnings (long panel). Table 8 shows the wage of each type relative to the overall average. For the short panel, the table further distinguishes wages in short-term (ST) and long-term (LT) jobs.[24]

It is clear from the table that the unstable types, High-U and High-N in the short panel data, and γ and δ in long panel data, on average have substantially lower wages or earnings. Further, earnings are particularly low for women who experienced a long non-employment spell (δs). The nature of the BEAM data do not reveal whether this is due to low wages or hours worked.

Nevertheless, these wage and earnings differences are small compared to overall wage or earnings inequality among the employed or that observed within types. This suggests that the labor market types we have described here differ more in their labor market flow rates than in productivity.

3.2.6 Summary

Comparing the two methods and types of data sources is not straightforward. Types are not identified in the same way, both due to differences in the methods (namely how the states are measured) and differences in the data (where the long vs short panel length plays a particularly important role, but also due to different sample characteristics and selection across countries). Still, they lead to conclusions that, while not identical, are broadly very similar, and hold across gender:

[24] We estimate wages by type and state directly from the C-LFS by matching measured wages conditional on LFS histories observed in the C-LFS. Our estimates are relative to the average wage in the population.

1. The labor market is characterized by substantial persistent heterogeneity in flows across states. Statistics capturing average labor market flows generally do not represent any of the types in the market well.

2. Most individuals, between 1/2 and 2/3 of the population, belong to a stable, strongly attached group. These individuals are employed most of the time, find jobs very quickly when non-employed, and earn higher wages.

3. A minority of the population is weakly attached, spends significant time non-employed, and earns lower wages when employed. In the following, we will refer to these groups – High-U and High-N in the short panel estimates, γs and δs in the long panels – as "unstable types".[25]

To us, the most glaring contrast between the methods is in flows between E and NE. The long panel method attributes the differences in time non-employed primarily to differences in separation rates rather than job finding rates. The short panel method finds the opposite.

[25] One could further separate these into a group with a high propensity for unemployment, and one with a high propensity to be out of the labor force. Ahn et al. (2023) find similar types and refer to them as "primary", "secondary" and "tertiary" segments.

Panel A: Short panels

Type	US (CPS, Hall-Kudlyak)					Canada (LFS)				
	Contribution to agg.			Flow rates (%)		Contribution to agg.			Flow rates (%)	
	E	U	N	E-NE	NE-E	E	U	N	E-NE	NE-E
Women:										
All-E	70	0	0	0	–	66	0	0	0	–
High-E	19	21	4	7.8	60.3	21	18	7	7.4	63.1
High-U	5	60	4	8.1	11.4	8	68	5	5.5	9.6
High-N	7	22	39	14.3	13.9	5	14	43	12.0	5.6
All-N	0	0	53	–	0	0	0	45	–	0
Single type									2.6	9.5
Men:										
All-E	77	0	0	0	–	72	0	0	0	–
High-E	16	25	7	7.2	65.5	16	22	8	7.4	61.9
High-U	4	59	5	9.3	12.4	7	67	7	6.6	11.2
High-N	4	16	35	11.8	8.0	5	11	44	8.5	9.0
All-N	0	0	54	–	0	0	0	41	–	0
Single type									2.1	14.4

Panel B. Long panel: Canada (BEAM)

Type	Women				Men			
	Contribution to agg.		Flow rates (%)		Contribution to agg.		Flow rates (%)	
	E	NE	E-NE	NE-E	E	NE	E-NE	NE-E
α	45	9	0.5	11.2	48	6	0.3	13.0
β	16	8	1.0	10.0	15	6	1.0	13.2
γ	22	22	2.7	12.6	24	30	3.3	15.3
δ	16	61	2.7	3.6	13	57	3.3	4.3
Total			1.3	7.1			1.4	8.8

Table 5: Contribution to aggregate employment and non-employment and flow rates by type

Note: Columns may not add up to 100 due to rounding.

Panel A: Short panel methods

Type	US (CPS, Hall-Kudlyak)				Canada (LFS)			
	Implied mean duration (months)				Implied mean duration (months)			
	Women		Men		Women		Men	
	E	NE	E	NE	E	NE	E	NE
All-E	+∞	na	+∞	na	+∞	na	+∞	na
High-E	12.7	1.6	14.0	1.6	13.6	1.6	13.5	1.6
High-U	13	9.9	10.7	8.1	18.5	10.4	15.2	8.9
High-N	7.2	14.1	8.8	11.9	8.4	17.8	11.8	11.1
All-N	na	+∞	na	+∞	na	+∞	na	+∞
Single type					38.6	10.6	47.7	7

Panel B: Long panel methods, women + men

B.1 Duration of employment spells

Type	US (LEHD, GMW)				Canada (BEAM)				
	Share of jobs that last				Share of jobs that last				Mean
	<1Q	1Q-4Q	5Q-8Q	>8Q	<1Q	1Q-4Q	5Q-8Q	>8Q	(months)
α	14	19	24	44	2	6	5	87	102
β	20	23	24	33	5	17	11	67	50.0
γ	36	32	19	12	12	47	14	27	20.0
δ					10	36	14	41	23.0
All					7	24	10	60	39.3

B.2 Duration of non-employment spells

Type	US (LEHD, GMW)				Canada (BEAM)				
	Share of spells that last				Share of spells that last				Mean
	<1Q	1Q-4Q	5Q-8Q	>8Q	<1Q	1Q-4Q	5Q-8Q	>8Q	(months)
α	79	16	5	0	1	97	1	0	8.1
β	39	56	5	0	41	19	39	0	8.6
γ	55	32	13	0	18	70	12	0	6.9
δ					9	33	12	46	23.5
All					16	57	14	13	12.4

Table 6: Duration of labor market states by type

Panel A: Duration of jobs

Type	Women					Men				
	Share of jobs that last				Mean	Share of jobs that last				Mean
	<1Q	1Q-4Q	5Q-8Q	>8Q	(months)	<1Q	1Q-4Q	5Q-8Q	>8Q	(months)
α	2	5	5	87	107	2	5	4	88	107
β	5	16	12	68	51	5	18	11	66	49
γ	10	45	15	30	23	14	48	13	25	18
δ	8	35	14	43	25	12	37	13	39	21
All	6	23	11	60	42	7	24	9	60	38

Panel B: Duration of non-employment spells

Type	Women					Men				
	Share of spells that last				Mean	Share of spells that last				Mean
	<1Q	1Q-4Q	5Q-8Q	>8Q	(months)	<1Q	1Q-4Q	5Q-8Q	>8Q	(months)
α	1	97	2	0	9	1	98	1	0	7
β	32	21	47	0	10	50	18	32	0	7
γ	15	70	15	0	8	20	70	10	0	6
δ	8	32	13	47	26	10	34	11	45	21
All	12	57	17	14	14	19	57	12	12	11

Table 7: Duration of labor market states by type and gender, Canada

Panel A: Short panel, wages: Canada (LFS)				
Type	Women		Men	
	ST	LT	ST	LT
All-E	–	1.10	–	1.09
High-E	0.94	0.88	0.86	0.9
High-U	0.65	0.71	0.69	0.8
High-N	0.72	0.73	0.75	0.72
All-N			–	–

Panel B: Long panel, earnings			
	US (LEHD, GMW)	Canada (BEAM)	
Type		Women	Men
α	1.21	1.16	1.13
β	0.81	1.04	1.10
γ	0.59	0.96	0.88
δ	–	0.76	0.82

Table 8: Wages or earnings by type, relative to the mean

4 The Costs of Labor Market Risk: A Simple Sketch

The Hidden Markov Models described and estimated above provide us with estimates of the population distribution across types ω and the transition matrices $\Pi(\omega)$ associated with each type. These generate a complete description of the dynamics of flows across labor force states in this model. However, the model on its own does not speak to how labor market risk experienced by agents is distributed in the population.

To describe how risk from labor market flows is distributed we need to specify the variation in earnings associated with transitions across states, and we need to calculate the willingness to pay for avoiding the volatility in earnings associated with these transitions. By focusing on earnings, we implicitly abstract from non-financial costs of job loss even as a growing literature reports sizable consequences of unemployment for mental health, subjective well-being, and even mortality.[26] We also assume that agents can not smooth consumption through savings or risk sharing within families, except by allowing for a replacement rate for lost earnings which we impose parametrically.[27]

What we do provide is a stylized calculation of the distribution of earnings risk based on estimates of the HK model that we obtained from the C-LFS and described in Section 3. We focus on the HK model since it makes a distinction between U and N, and between short- and long-term jobs, both of which are relevant to our calculations. The estimated transition matrices $\Pi(\omega)$ provide us with projected labor

[26] Helliwell & Huang (2011), Winkelmann & Winkelmann (1998), Eliason & Storrie (2009), and Sullivan & von Wachter (2009).

[27] See Lentz & Tranaes (2005), Low et al. (2010) and Altonji et al. (2013).

flows associated with each type ω. They also imply ergodic distributions $\mathbf{n}^*(\omega)$ with elements $n_k(\omega)$ capturing the steady state proportion of a type in a given state $k \in 1, ..., K$. For each type, we also have estimates of the relative wages associated with short-term and long-term employment in Table 8. For example, a high-U individual in long-term employment is assumed to earn approximately 1.16=0.80/0.69 as much as in short-term employment.

We assume infinitely-lived risk-averse individuals with flow utility $u(c) = c^{1-\sigma}/(1-\sigma)$ who discount the future at rate β. σ is the coefficient of relative risk aversion. Since these agents do not have access to a savings technology, they consume their current income in each period. When employed, income is given by the earnings provided by the estimates in Table 8. When unemployed or out of the labor force, we assume they receive income b. This income captures benefits and transfers, plus possibly the value of leisure, converted into consumption units.

We calculate the constant income stream $\bar{w}_k(\omega)$ that generates the same utility as that provided by the fluctuating income stream implied by the transition model. This approach follows Lucas (1987). Low et al. (2010) employ a similar criterion in a life cycle model. This constant income depends on the type of the individual and the current state. Because agents are risk averse, this constant income stream is smaller than the average income the individual in state k will receive over the remainder of her life-time, which we call permanent income. The difference between the constant income stream and permanent income is a risk premium. We measure this risk premium as a percentage of permanent income and denote it by $\lambda_k(\omega)$. This is the fraction of permanent income that the individual is willing to pay to avoid labor risk.

We can also calculate this risk premium under the veil of ignorance, assuming that individuals know their type, but not the labor force state once the veil has been lifted. Denote this by $\lambda(\omega)$.

Besides this measure of the willingness to pay to avoid the income volatility associated with flows across the different states, we also calculate an alternative measure of the percent of income each type would be willing to forego to avoid a single episode of losing employment in a long-term job and transitioning to unemployment. We thus compute $\mu(\omega)$, the permanent percent reduction in income on the long-term job that reduces welfare by as much as the move to unemployment.

Panel A of Table 9 shows the cost μ of losing a long-term job and entering unemployment. Panel B shows the willingness-to-pay to avoid income volatility λ. Because All-E and All-N types do not experience job loss, we can't calculate μ for them, and λ will trivially equal zero 0. When we report λ's, we report both the $\lambda(\omega)$ obtained under the assumption that agents are behind the veil of ignorance and the range (in brackets) for the state-specific $\lambda_k(\omega)$'s. For both panels, we also report the weighted average across types, as well as the values obtained using the estimates for the single-type economy.

Both μ's and λ's depend on the replacement rate b for which we consider two values, 0.4 and 0.65, frequently used in the literature.[28] They also depend on the coefficient of relative risk aversion σ, which we set equal to 1.5. We assume that agents discount the future at an annual rate of 4%.

Panel A shows that the cost of job loss differs markedly across labor market types. For example, if the replacement rate is 0.4, then men of the High-N type employed

[28] 0.4 is the typical replacement rate used in quantitative DMP models applied to the United States (see e.g. Shimer 2005). Adding to this the value of additional leisure, Hall & Milgrom (2008) use a value of 0.65. We impose a common replacement rate across types, which we apply to the type-specific mean wage in a long-term job shown in Table 8. A plausible alternative is for b/w to differ by type.

in long-term jobs are willing to give up 5.5% of their income to avoid a transition from employment in a long-term job to unemployment. The High-E types are only willing to forego 2.2% of their income to avoid the same transition. These High-E individuals have much higher job finding rates, so periods in non-employment are less costly for them. High-N women are willing to give up 10.2% of their incomes due to the very low transition rates to employment that they face from the non-employment states.[29]

It is striking how much the cost of fluctuations λ differ by labor market type. First, by construction, there is no cost for the All-E and All-N types who always occupy the same state.[30]

Second, the cost of fluctuations for the high-U and high-N types far exceeds that experienced by other types. With a common replacement rate b, the cost almost entirely depends on the mean rate of non-employment for the type, which is much higher for the high-U and high-N types. The average cost of fluctuations across types is greater for women because more women belong to the high-N type.

In addition, there are two more subtle but important results. First, the average cost of fluctuations across types is far lower than the cost that would be inferred ignoring types (last column). This is because the latter conflates heterogeneity across types with risk. Second, there is little variation in willingness to pay to avoid income risk across states within types. This is because at standard discount rates, the flow rates are high enough for all types to generate rapid conversion back to the ergodic distribution from any starting point.

[29] The High-N types frequently transition to N after a transition to U. Thus, the low job finding rates out of N are particularly important for them.

[30] Is it plausible that a large part of the population faces no losses due fluctuations? In short panels, it is inherently difficult to identify the probability of rare events, like job loss of high-tenure workers. We return to the question of the cost of job loss for stable workers in Section 7.3.

	Type						Single
	All-E	All-N	High-E	High-U	High-N	Average	type
Panel A: The cost of job loss μ							
				Women:			
$b = 0.4$:	–	–	1.0	4.5	10.2	4.6	4.4
$b = 0.65$:	–	–	0.4	1.9	4.4	2.0	1.9
				Men:			
$b = 0.4$:	–	–	2.2	4.0	5.5	3.5	3.0
$b = 0.65$:	–	–	1.2	1.7	2.3	1.6	1.3
Panel B: The cost of labor market risk λ							
				Women:			
$b = 0.4$:							
veil of ignorance	0	0	5.4	12.7	13.2	3.8	9.0
by state	[0,0]	[0,0]	[5.4,5.5]	[12.5,12.9]	[13.1,13.5]	[3.8,3.8]	[8.8,9.7]
$b = 0.65$:							
veil of ignorance	0	0	1.4	3.0	3.0	0.9	2.2
by state	[0,0]	[0,0]	[1.4,1.4]	[3.0,3.1]	[3.0,3.1]	[0.9,0.9]	[2.1,2.3]
				Men:			
$b = 0.4$:							
veil of ignorance	0	0	5.1	12.7	14.3	3.3	6.2
by state	[0,0]	[0,0]	[5.0,5.2]	[12.6,12.9]	[14.2,14.3]	[3.2,3.3]	[6.0,7.0]
$b = 0.65$:							
veil of ignorance	0	0	1.2	3.0	3.5	0.8	1.6
by state	[0,0]	[0,0]	[1.2,1.2]	[3.0,3.1]	[3.4,3.5]	[0.8,0.8]	[1.5,1.8]

Table 9: The costs of labor market risk

Data: CLFS, 2001-2019, respondents aged 25-54.

5 Heterogeneity and Duration Dependence in Job Finding and Loss Rates

We now relate the main findings from the short and the long panel approaches to the literature on duration dependence, starting with the job finding rates, and then considering the job separation rates.

5.1 Duration dependence in job-finding

Morchio (2020) and Ahn & Hamilton (2020) are but two recent examples of a vast literature on heterogeneity in experienced unemployment. Morchio (2020) reports that 10% of respondents to the National Longitudinal Survey of Youth 1979 account for two-thirds of prime-age unemployment, whereas 58% do not experience any unemployment at all. This heterogeneity in experienced unemployment suggests that individuals differ in the rate at which they lose and find jobs. Ahn & Hamilton (2020) find that about 75% of newly unemployed in the CPS belong to a type likely to exit unemployment rapidly, while the rest is much less likely to exit quickly. They, however, also find evidence of genuine duration dependence (GDD) in job finding rates, as agents in their model transition from short- to long-term unemployment with lower job finding rates.

Separating and quantifying how much heterogeneity in job finding rates as opposed to GDD contributes to heterogeneity in experienced unemployment is challenging. Even without any GDD, observed job finding rates decline as durations lengthen and those that remain unemployed increasingly consist of those with the lowest exit rates. Heckman & Singer (1984*a*) notes the result: "Uncontrolled unobservables bias estimated hazards towards negative duration dependence."[31] Single-spell data does not allow separating heterogeneity from GDD without imposing additional structure. Lancaster (1979), in his early work on this topic, is but a forerunner of a large literature imposing proportional hazard assumptions. Parametric restrictions that permit identification such as the proportional hazard assumption are, however, typically ad-hoc and do not follow from economic models. As Machin & Manning (1999) put it: "it does not really seem possible in practice to identify separately the effect of heterogeneity from that of duration dependence without making some very strong assumptions about functional forms which have no foundation in any economic theory."[32]

In this section, we connect the new literature on heterogeneity in labor force transitions reviewed in prior sections to the work on separating dynamic selection and GDD. Before we do so, we (selectively) review some recent attempts to approach this problem. We refer the reader interested in details to the foundational papers by Heckman & Singer (1984*c*,*a*,*b*), Lancaster (1979) and Honoré (1993). Van den Berg (2001) reviews the specification and identification of duration models with a particular focus on proportional hazard models. A particularly useful starting point for the novice is Section 5 of Machin & Manning (1999).

[31] Knowledge of the direction of the bias allows for non-parametric tests of negative GDD. However, hazards typically exhibit negative observed duration dependence, so these procedures are of little help in this context.

[32] Machin & Manning (1999), p. 3111

5.1.1 GDD or Dynamic Selection in Job Finding Rates? A Selective Literature Review

Resume audit studies rely on interview call back rates and their variation with spell durations implicit from CVs for this purpose (Eriksson & Rooth (2014), Ghayad (2014), Kroft et al. (2013), Farber et al. (2016), Farber et al. (2019), Nüß (2018)). Even though these studies rely on similar methodologies, they do not all arrive at similar results, possibly because they differ in the populations they study. (Eriksson & Rooth (2014), Kroft et al. (2013), and Ghayad (2014) find evidence for GDD in call back rates; Farber et al. (2019), Nüß (2018) report evidence only for long spells but not at short duration; Farber et al. (2016) and Nunley et al. (2017) find no evidence at all. Furthermore, interview callbacks are not equivalent to job offers, and Jarosch & Pilossoph (2019) use a calibrated frictional labor search model to argue that observed effects of durations on callback do not translate to a significant decline in job offers. On balance, this literature does provide some support for GDD in call-back rates, particularly among younger and less educated workers. However, the literature has not converged on an answer to the question of whether and how callback rates and possibly job finding rates evolve with duration.

Simple calibrated models of matching allowing for duration dependence in job finding rates from unemployment (Kroft et al. (2016)) and also from non-participation (Kroft et al. (2019)) have been successful in matching the dynamics of the US and the Canadian labor market over the Great recession. More recently, similar approaches had success in matching labor market dynamics during the pandemic recession (Gallant et al. (2020), Forsythe et al. (2022)) once recall unemployment (see Section 6) was accounted for. Kroft et al. (2016) report little role for heterogeneity in job finding rates by observable demographic characteristics. None of these papers do directly speak to the question of heterogeneity in unobservables.

Still other papers rely on functional form assumptions. For example, Hornstein (2012) proposes a model of unemployment with two latent states (short and long-term unemployment) that differ in the outflow rates from unemployment.[33] He reports that the evolution of the unemployment duration distribution is well described by significant differences in these outflow rates, with relatively little room for duration dependence arising from switching between the latent states. Another approach is to impose restrictions on the variation in job finding rates with observables. For example, Van den Berg & Van Ours (1996) estimate the extent of heterogeneity and duration dependence using variation in aggregate outflow rates across time. They likewise find an important role for heterogeneity in explaining duration dependence. Machin & Manning (1999) summarize the literature on duration dependence in Europe up until the end of the 20th century. With the exception fo the UK, they find little evidence for true negative duration dependence in exit rates from unemployment (including to non-participation) once one controls for observable heterogeneity.

Multiple spell data promises a way forward if the unobserved heterogeneity in exit rates from unemployment is constant across spells (Honoré (1993)). Intuitively, repeated unemployment spells of the same individuals allow controlling for permanent heterogeneity in outflow rates and thus allow identifying the causal effect of durations on job finding rates. Once this has been identified, it is possible to identify the distribution of heterogeneity in job finding rates in the population. Alvarez et al. (2023) for example estimate a model with GDD and unobserved heterogeneity on individuals with least two non-employment spells in social security data from Austria. Mueller & Spinnewijn (2023) use machine learning tools on Swedish administrative

[33] This is in line with evidence from for example Krueger et al. (2014).

data on recipients of unemployment benefits with both repeated spells and an unusually rich vector of observables to identify the contribution of dynamic selection and GDD to observed duration dependence in job finding. They find that dynamic selection can account for 49-88% of observed duration dependence in their data.

On balance, our (incomplete) reading of the literature suggests that there is indeed heterogeneity in job finding rates, but that it is not clear whether there is GDD in job finding rates.

5.1.2 New Evidence from Short and Long Panels

The papers reviewed in Table 2 all identify heterogeneity in transition rates using data that includes multiple spells. By allowing for labor force states such as short-term and long-term unemployment, they can also accommodate GDD in job finding rates. Ahn et al. (2023) in fact do precisely this and allow for short- and long-term unemployment as latent states of their model. By contrast, Hall & Kudlyak (2022) do not allow for short- and long-term unemployment as distinct labor force states. All GDD in job finding rates among the non-employed in their model arises purely from transitions between U and N. All observed duration dependence conditional on labor force state U or N in job finding rates in their model thus necessarily stems from dynamic selection on unobserved types.

When considering job separation rates, Ahn et al. (2023) and Hall & Kudlyak (2022) again take very different approaches. Hall & Kudlyak (2022) do allow for short-term and long-term employment as structural latent states that are not directly observed, which allows them to match genuine duration dependence in job loss rates. By contrast, Ahn et al. (2023) do not allow for multiple structural latent states conditional on employment. Their model thus attributes all observed duration dependence in job loss rates to dynamic selection.

These differences in their structural assumptions mean that the two papers come to very different conclusions about the presence and type of GDD in transition rates in data from the Current Population Survey.

Short panel methods using data such as the CPS or in our case the C-LFS will typically only include a few, censored non-employment spells. This limits the ability of these models to distinguish and allow for more than a small number of latent states and types. We suspect that this limitation of the short panels in not generating enough repeated spells results in the lack of agreement between Hall & Kudlyak (2022) and Ahn et al. (2023).

By contrast, the long panel approach favored by Gregory et al. (2021) identifies type heterogeneity using many years of observations (1997-2014 in their case). It is thus natural to ask what we can learn about duration dependence conditional on types identified in the data from this approach. This approach does not allow looking specifically at exit or job finding among the unemployed but rather pools both the unemployed and those not-in-the-labor force.

Figure 1 displays the observed Non-employment to Employment hazards for males and females in the BEAM data as a function of the duration of non-employment up to 60 months. Surprisingly, these observed transition rates do not universally decline with observed non-employment durations. Rather, they increase and peak between 8 and 12 months of non-employment. Given that heterogeneity implies negative bias in observed job finding rates, this is by itself strong evidence of positive GDD in job finding rates through the first year of non-employment. This pattern, however, is not generated by traditional job finding as conceived of in standard job search models.

Rather, the pattern reflects recall non-employment, which we discuss in more detail in Section 6 below. Recall non-employment spells end when the separated individual rejoins their former employer.[34] This includes temporary layoffs in response to low demand, non-employment spells for parental and other leave-taking, and seasonal leave.

[Figure 1 here with caption: Non-employment to Employment Transition Rates by Gender and Duration of Non-employment Spell]

Figure 2 recodes all recall non-employment spells as employment and displays the non-employment to employment hazards (the job finding rate) for the remaining spells. This is equivalent to the job finding rate for those who *ex post* are not recalled. Here, we observe the expected decline in job finding rates throughout the first 2 years of non-employment, first rapidly and then somewhat slower. The vast majority of these spells is shorter than 2 years, with only 13% of spells lasting longer than 2 years (see Table 7). After 2 years, job finding rates stabilize and oscillate around their long-run levels through the remainder of the duration distribution.

[Figure 2 Here with caption: Job Finding Rates by Gender and Duration excl. Recall.]

As alluded to above, data from multiple spells allow separating heterogeneity and dynamic selection from GDD. The clustering approach following Gregory et al. (2021) relies on long panels with repeat spells to identify different types of agents. The obvious next step then is to ask whether there is evidence for GDD conditional on the type classification imposed by this method. Figure 3 plots the job finding

[34] In labor Force Surveys such as the CPS and the C-LFS, recall is typically measured prospectively by asking respondents whether they expect to be rehired. The C-LFS reports expected recall for both U and N, while the CPS includes this information only for U (Kroft et al. (2019)). In administrative data, these spells can be accurately coded retrospectively using the employer information.

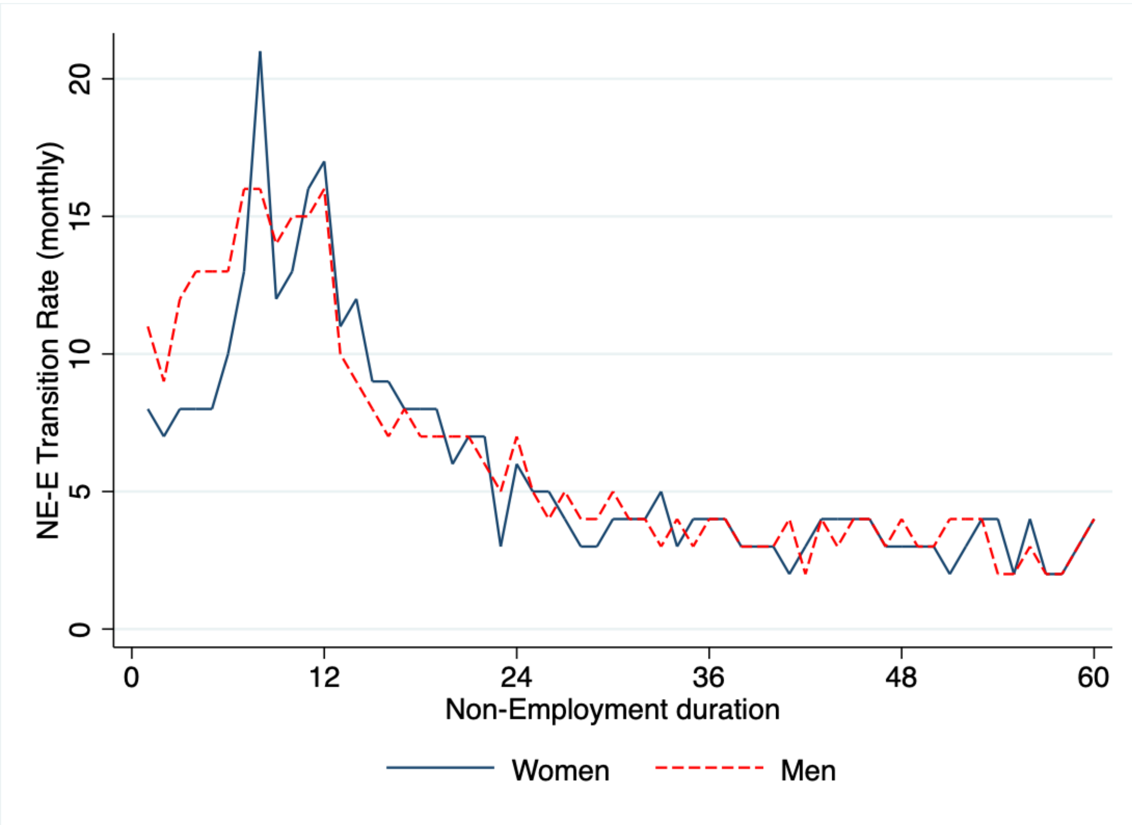

Figure 1: Non-employment to Employment Transition Rates by Gender and Duration of Non-employment Spell

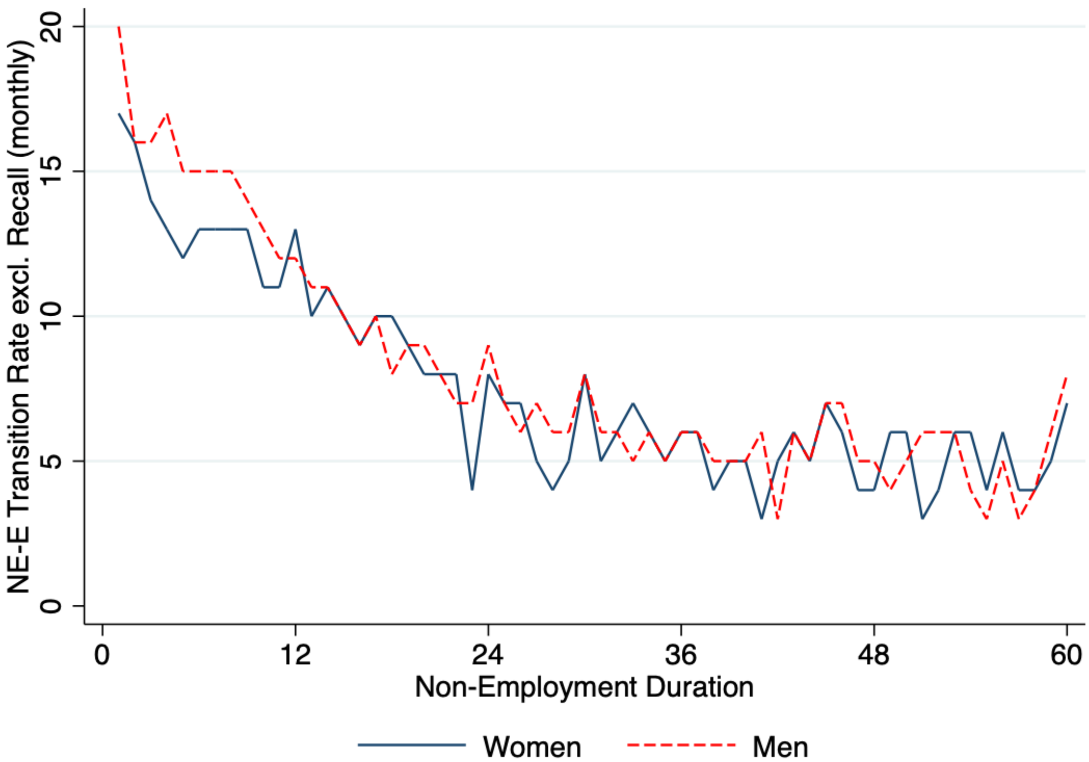

Figure 2: Job Finding Rates by Gender and Duration excl. Recall

rate excluding recall spells against the spell duration conditional on both gender and type.[35] We observe for both males and females evidence that job finding rates decline with duration. For both genders and most types, we also see a spike in hazards around 12 months, suggesting at least some non-employment spells that are systematically timed to last a calendar year. Such spells that last precisely a year might for example be due to parental leave, which individuals are entitled to for a year. Broadly though, we see declining job finding rates across all types. A rough comparison with Figure 2 indeed suggests that job finding rates decline at roughly the same rate within types as in the aggregate.

[Figure 3.1 and 3.2 here: Horizontally in a row. Below 3.1 (a) Women Below 3.2 (b) Men Caption: Job Finding Rates by Gender and Duration and Type excl. Recall]

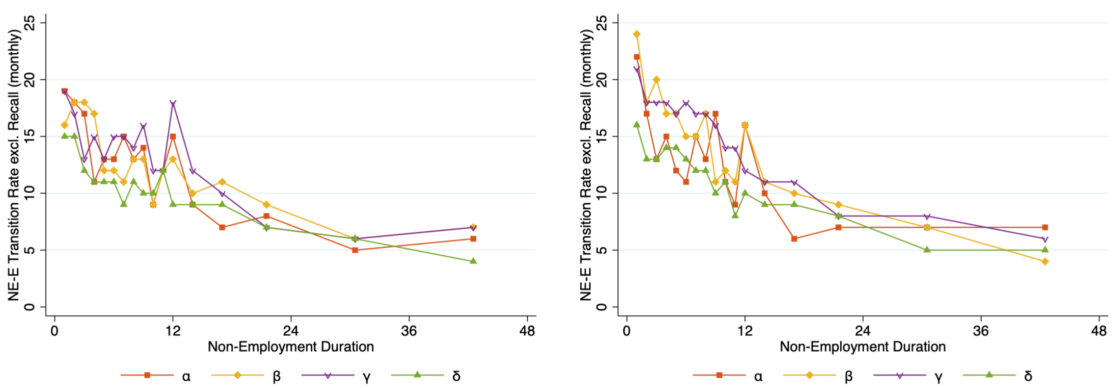

(a) Women

(b) Men

Figure 3: Job Finding Rates by Gender and Duration and Type excl. Recall

[35]Due to small sample sizes and disclosure requirements from Statistic Canada, some duration categories above 12 months have been binned into larger categories. In this and following figures, the duration mid-point of each category is shown. Similarly, we occasionally need to truncate hazards at an earlier point in the following graphs to meet disclosure requirements.

The similarity of job finding hazards for the different types shown in Figure 3 suggests a limited role for dynamic selection on unobservables in driving negative duration dependence in the aggregate job finding hazard. We investigate the role of types further in Figure 4. This figure shows the aggregate job finding hazard (excluding recall) as shown in Figure 2 above together with a counterfactual hazard that is a fixed-weight average of the type-specific hazard functions. (The aggregate hazard by construction is a similar average, but with weights that shift reflecting the type composition of the non-employment pool with duration.) We use the share of non-employment accounted for by the different types as weights (see Table 5 above).

Because the counterfactual hazard is computed with constant weights, and the aggregate hazard allows for shifting weights, the difference between the two shows the contribution of dynamic selection on unobservables to duration dependence. The lack of a difference shown in the figure then suggests that the declining job finding hazard largely reflects GDD, with little or no contribution of dynamic selection. This finding reflects the relatively small differences in job finding rates across types reported in Panel B of Table 5.

[Figure 4.1 and 4.2 here: Horizontally in a row. Below 4.1 (a) Women Below 4.2 (b) Men Caption: Job Finding Rates by Gender and Duration and Type excl. Recall: The Role of Types]

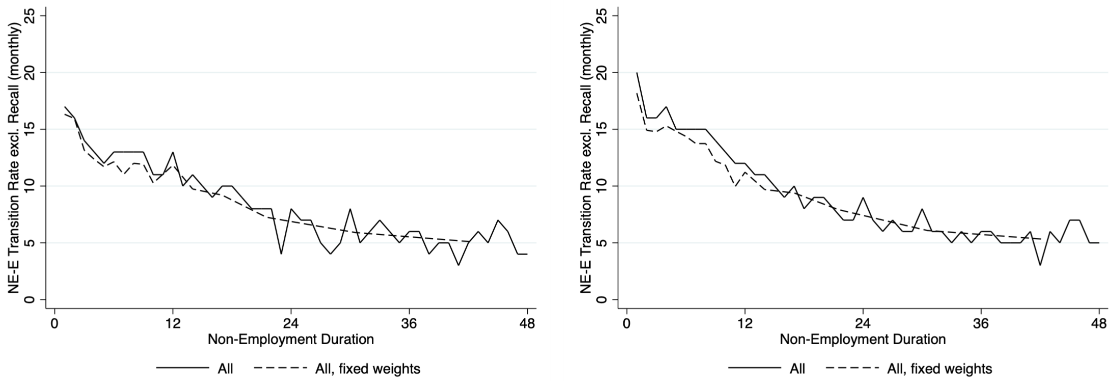

(a) Women

(b) Men

Figure 4: Job Finding Rates excl. Recall: The Role of Types

5.2 Duration dependence in separations

We can also explore how large the role of unobserved heterogeneity is in explaining outflow rates from employment. Figure 5 shows the observed employment separation rates by gender and duration which follow a distinct hump-shaped pattern, as implied by the match-learning model by Jovanovic (1979). During the first few months, separation rates increase before declining sharply after about 6-9 months of job tenure.

Figure 6 then demonstrates the same pattern within types. Differently from the job finding rates, we do observe large difference in job loss rates between the labor force types, especially in the first few months.[36]

[36]Some cross-country differences analogous to these type differences have been documented. Donovan et al. (2023) show that separation rates are systematically larger and separation hazards steeper in poorer countries, and interpret this as reflecting selection on match quality. Poschke (2023) analyzes these cross-country patterns using a search and matching model with latent states and attributes them to a low rate of screening upon hiring combined with frequent shocks to match productivity.

We then use the same fixed-weight as before to evaluate in Figure 7 how important dynamic selection and GDD are in generating the observed duration patterns in job separations. As with the job finding rates, this approach also suggests that indeed durations generate genuine negative duration dependence, at least following the first few months. However, there is some evidence that dynamic selection plays an important role, too.

[Figure 5 here Caption: Separation Rates by Gender]

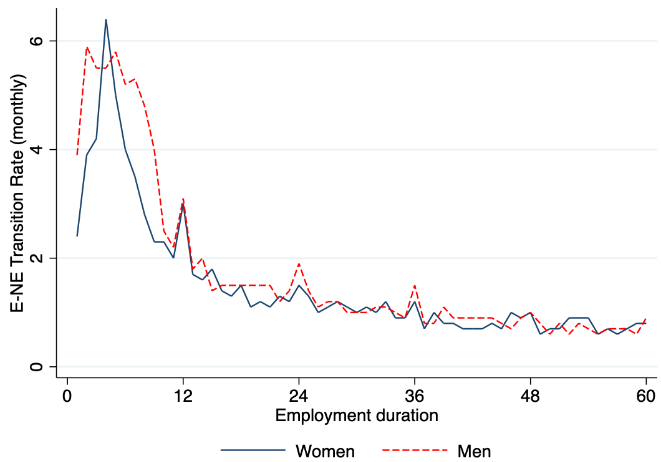

Figure 5: Separation Rates by Gender

[Figure 6.1 and 6.2 here: Horizontally in a row. Below 6.1 (a) Women Below 6.2 (b) Men Caption: Separation Rates by Gender, Duration and Type]

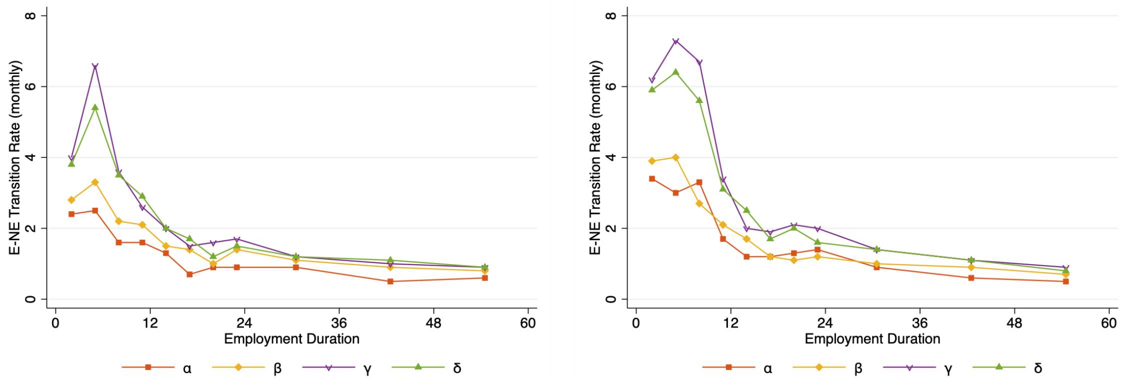

(a) Women (b) Men

Figure 6: Separation Rates by Gender, Duration and Type

[Figure 7.1 and 7.2 here: Horizontally in a row. Below 7.1 (a) Women Below 7.2 (b) Men Caption: Separation Rates: The Role of Types]

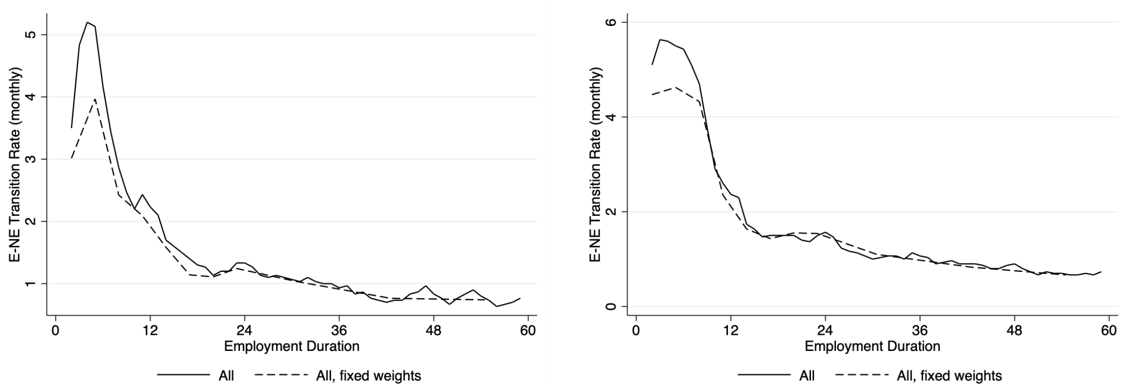

(a) Women (b) Men

Figure 7: Separation Rates: The Role of Types

We thus find a role for dynamic selection on unobservables in shaping employment separation rates, but less so for job finding rates. The latter finding contrasts with conclusions of other authors, such as Hornstein (2012), Alvarez et al. (2023), and Ahn (2023), who find that dynamic selection explains a lot of the duration dependence in job finding rates. We believe that this discrepancy is, at least in part, due to differences in the transitions used to identify heterogeneity. As emphasized above, we exploit a long panel which allows classifying individuals based on dynamics observed across their careers with repeated spells of employment and non-employment. Moreover, we also force types to identify heterogeneity in both job separations and job finding rates. It is therefore possible that the discrepancy arises because the above-cited authors rely on moments related to the duration structure of unemployment only. This might well result in them finding more heterogeneity by unobservables in job finding rates than we do.

Using the Canadian BEAM data, and following the clustering approach, we do find that labor force types differ significantly more in separation than in job finding rates, as shown in Table 5 above. As a result, we attribute less importance to type heterogeneity in job finding rates. How types estimated on specific transition rates align with those estimated on the full set of transitions remains a topic to be explored in more detail.

6 Recall Unemployment and Non-employment

Recent literature has also drawn attention to the importance of recall, i.e. a return to the previous employer after an unemployment or non-employment spell. For example, Fujita & Moscarini (2017) document that in data from the US Survey of Income and Program Participation (SIPP), over 40 percent of workers return to their previous employer after a separation into unemployment. The recall rate is still around 20 percent for workers who separated permanently from their employer, and did not expect to be recalled. It is much larger for those laid off temporarily. Fujita & Moscarini (2017) also find that recalled workers had a longer previous employment spell, and experience a shorter unemployment spell. In their data, recall is the main component driving duration dependence of unemployment exit. These findings build on and extend earlier results using US surveys by Katz (1986) and Katz & Meyer (1990), showing large levels of recall and the importance of recall for the shape of the job finding hazard.[37]

Large rates of recall have also been documented in data from other countries. Jansson (2002) finds that in Sweden, 45 percent of completed unemployment spells result in recall. Alba-Ramírez et al. (2007) find this number to exceed a third in Spanish Social Security data, and Nekoei & Weber (2015) document an overall recall rate of 35 percent in Austrian administrative data. While recall rates are higher after temporary layoffs, they are nevertheless significant after permanent separations, when recall was not expected. This is important, since the latter account for the majority of separations. These authors also show that recall matters beyond seasonal industries or manufacturing.[38]

A core question in the literature concerns the shapes of the recall and new job finding hazard, and which of the two accounts for the typical negative duration dependence in job finding. To this end, authors typically estimate flexible duration models, often controlling for both observed and unobserved heterogeneity. The latter otherwise induces negative duration dependence in hazards, as also mentioned above. Examples of this approach are Fallick & Ryu (2007) and Alba-Ramírez et al. (2007), who model unobserved heterogeneity using a discrete distribution, following

[37] Further classic references on the importance of recall are Feldstein (1975) and Lilien (1980).

[38] The data on which we rely precedes the COVID pandemic during which recall non-employment reached unprecedented heights. For a discussion of the role of recall non-employment during that time, consider Gallant et al. (2020) and Forsythe et al. (2022).

the recommendation of Heckman & Singer (1984*c*).

Findings on the shape of hazards differ across countries and data sources. Fallick & Ryu (2007) and Alba-Ramírez et al. (2007) find negative duration dependence only in the recall hazard in US survey data from the Panel Study of Income Dynamics and Spanish Social Security data, respectively. Fujita & Moscarini (2017) similarly find much stronger negative duration dependence in recalls than in new job finding over the first 6 months of an unemployment spell in US SIPP data. Nekoei & Weber (2015), in contrast, find negative duration dependence of new job finding and a hump-shaped recall hazard in Austrian unemployment insurance records.

These approaches allow for and estimate separate unobserved heterogeneity specifically in recall and job finding rates. In this section, we explore recall rates across the broader labor market types estimated above.

Figure 8 shows the hazards of recall (dashed, blue), new-job finding (dash-dot, red), and the combined job finding hazard (solid, black). The recall hazard first increases, peaks around eight to twelve months, and then decreases. After about a year, it quickly declines towards zero, dropping below one percent around two years, and then slowly declining further. The overall job finding hazard reflects both the new job and the recall hazard. At long non-employment durations, the probability of recall is much smaller than that of finding a new job, and the overall hazard reflects almost exclusively the latter. But at shorter durations, the overall hazard clearly inherits the hump shape of the recall hazard. In fact, the probability of recall exceeds that of finding a new job from non-employment during months five to twelve (men) or fourteen (women). Moreover, while there is negative duration dependence in the new job finding rate, most of the decline in the overall job finding rate after the first year of non-employment reflects the negative duration dependence in that part of the recall hazard.

The plateau at eight to twelve months in the recall hazard for men lines up closely with the maximum duration of unemployment insurance benefits in Canada. The peaks of the recall hazard for women (at eight, eleven and twelve months) are even more pronounced. This likely reflects the fact that we are measuring recall out of *non-employment*, not *unemployment* (which cannot be distinguished in the BEAM data), combined with regulations affecting leaves in Canada. For example, Canada has generous parental leave policies that allow a mother who combines maternity and parental leave to take a leave of about a year at a 55% replacement rate. (Fathers also have access to benefits, but typically take them up for a shorter period.) Data on the reason for non-employment will allow to pin down the contribution of different types of separations and leaves to recall more precisely, beyond the broad distinction between temporary layoffs and permanent separations that has already been explored in the literature (see e.g. Fujita & Moscarini 2017, Nekoei & Weber 2015). These findings also suggest substantial differences between recall out of unemployment and recall out of non-employment. Recall out of unemployment peaks much earlier and has been the focus of much of the literature. Recall out of non-employment includes recall from those not in the labor force, a group which accounts for much of the non-employment to employment transitions. Given their quantitative importance, recall for this group (see e.g. Fujita & Moscarini 2017) warrants further investigation.

[Figure 8.1 and 8,2 here: Horizontally in a row. Below 8.1 (a) Women Below 8.2 (b) Men Caption: The Probability of Recall by Gender and Duration of Non-Employment]

Table 10 shows the average probability of recall for the non-employed, and the share of non-employment spells that end in recall. Slightly more than half of all non-

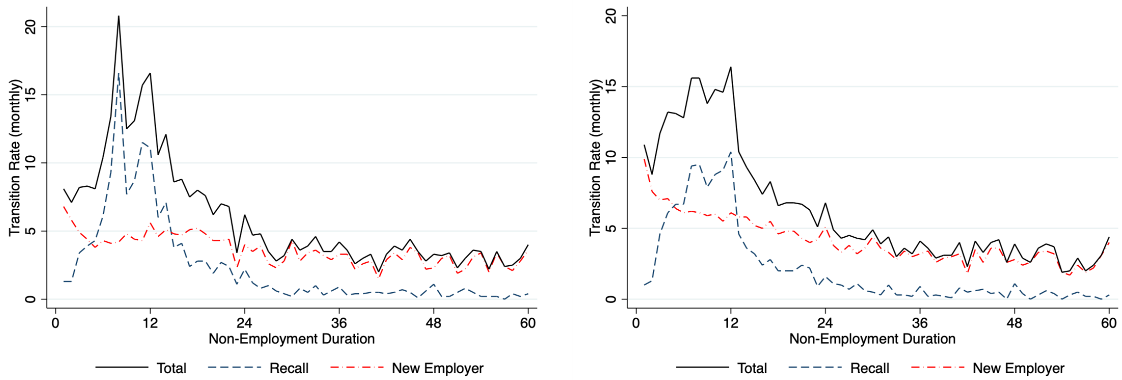

(a) Women (b) Men

Figure 8: The Probability of Recall by Gender and Duration of Non-Employment

employment spells end in recall. This fraction is about ten percentage points larger for women than for men. These high rates of recall are in line with the literature discussed above.[39]

The table also shows how recall rates differ by labor market type. It is clear that the incidence of recall is high for all types. Among women, the share of non-employment spells that end in recall is lowest among δs, for whom it is about ten percentage points lower than for αs or γs, but exceeds one half for all types. Among women, the share is roughly ten percentage points larger among αs than among the other groups, but it comes close to one half for all types. Variation in the average probability of recall is similar. It is only about half as large for δs as for the other types. The share of non-employment spells that end in recall is nonetheless large even for δs because their rate of finding new jobs is also lower.

The recall hazards shown in Figure 9 are strikingly similar across types, apart from the generally lower level for δs. The peaks of the hazard vary slightly by type,

[39]This figure is similar in Canadian administrative data for 1984, see Robertson (1989).

with the peak at eleven and twelve months dominating for αs, and the peak at eight months dominating for the other three types. This suggests that different types experience non-employment episodes for different reasons – a promising avenue for further research.

	Share of NE spells that end in recall			Probability of recall		
Type	All	Women	Men	All	Women	Men
α	60.3	64.2	56.0	5.2	5.4	4.9
β	52.0	57.7	47.2	4.5	4.8	4.2
γ	52.3	60.1	47.8	5.2	5.3	5.2
δ	49.6	52.6	46.8	2.5	2.7	2.4
Total	53.1	58.4	49.1	4.4	4.5	4.3

Table 10: The probability of recall

[Figure 9 here Caption: The Probability of Recall by Type and Duration of Non-Employment]

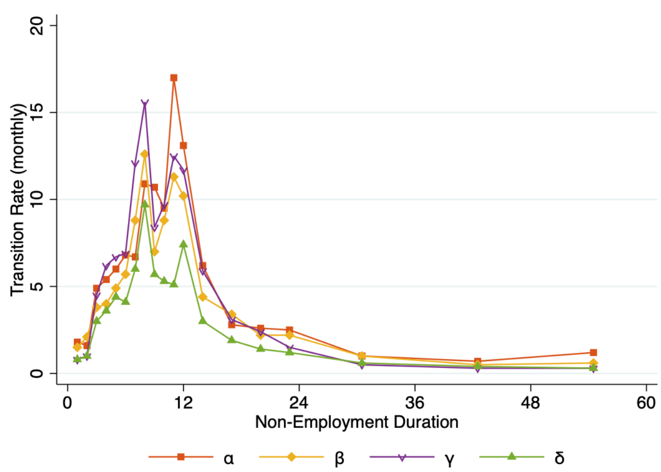

Figure 9: The Probability of Recall by Type and Duration of Non-Employment

7 Enriched employment states: job ladders

Following Burdett & Mortensen (1998), a large literature has emerged featuring models with search frictions and heterogeneous jobs which generate a "job ladder".[40] Such a ladder arises when workers share a common ranking of vertically differentiated jobs, but search frictions prevent them from directly accessing their preferred job out of non-employment. They thus accept a satisfactory job and continue their search while on the job, moving on to better jobs when they find them. This generates a career profile: workers move up the job ladder via job-to-job transitions.[41] Rungs of the job ladder resemble labor force states in that they condition the rate at which individuals transition across rungs as well as between employment and non-employment. The process of moving up the job ladder may eventually be disrupted by a layoff shock. For those at the top of the ladder, the costs of being knocked off the ladder and losing this valuable perch might be long lasting and sizeable even if they are reemployed quickly. In this section, we discuss recent evidence on job-to-job flows, job ladders, and the costs of job loss, and provide additional evidence from a perspective of heterogeneity across labor market types as measured above.

In empirical work, the first problem that arises is to construct the job ladder. Given the difficulty of separating job and match attributes, the literature has typically identified jobs with firms. Differences in value or economic rents associated with the rungs of the ladder can arise when employers are homogeneous but have heterogeneous wage policies as in Burdett & Mortensen (1998). However, most job ladder models assume that jobs are differentiated by productivity.[42] Unfortunately,

[40] See Moscarini & Postel-Vinay (2018) for a recent review. Seminal references include Postel-Vinay & Robin (2002) and Cahuc et al. (2006).

[41] For direct evidence on search by rung of the ladder, see Faberman et al. (2022).

[42] Besides compensation, jobs can also differ along other dimensions, making them more or less

productivity measures are rare in data with information on job flows forcing researchers to proxy for firm productivity using variables such as firm size (see, for instance, Bartelsman et al. 2013). Alternatively, job ladders can be constructed using differences in firm's average wages or earnings. These are typically available in matched employer-employee data and are known to correlate with its size (see e.g. Brown & Medoff 1989) and productivity (Abowd et al. 1999).

Once the ladder has been constructed using either firm size or compensation measures, the literature then looks for evidence on the ladder using job-to-job flows. Job ladders embed the idea that jobs higher up the ladder are able to attract workers and tend not to lose employees to lower-ranked firms. Evidence on ladders centered on the ability to attract workers thus comes in the form of showing that *gross poaching rates,* defined as the fraction of a firm's total hires that come from other firms rather than from non-employment, increase across firm rungs (see e.g. Moscarini & Postel-Vinay 2008). Alternatively, Haltiwanger et al. (2018) propose using the *net poaching rate,* which includes the outflow rate to other firms. This measure thus incorporates information related to both the ability to attract and to retain workers.

Among researchers constructing ladders based on firm size, Moscarini & Postel-Vinay (2008, 2016) use the US Survey of Income and Program Participation (SIPP) to document that the gross poaching rate increases in firm size. Among those relying on compensation measures to construct job ladders, Haltiwanger et al. (2018) show evidence for a wage ladder in US LEHD data, while they fail to find evidence for a

attractive to agents. These might include job security (Jarosch 2023), health insurance (Dey & Flinn 2005) or other job amenities (Hall & Mueller 2018, Albrecht et al. 2018, Taber & Vejlin 2020). When the production technology induces sorting of workers and firms – as is the case in a number of recent contributions to the search literature – workers may rank jobs in different ways, and thus ladders will be specific to worker type. See also Lise & Robin (2017) or Borovičková & Macaluso (2024).

size based ladder.[43,44]

Recent work has also directly classified firms based on worker flows. Using Danish matched employer-employee data, Bagger & Lentz (2019) document a stable ladder in terms of firms' gross poaching rank, which they find to be correlated with a firm's wage rank. Sorkin (2018) rely on revealed preferences as manifested in job-to-job flows to estimate a ladder in terms of net poaching in US LEHD data. This ranking can deviate from a wage-based ranking in that firms are also differentiated by amenities.[45] Overall, the evidence based on firm size, compensation, or revealed preferences does suggest that there is an operative job ladder.

Two of the implications of an operative ladder stand out in the context of this chapter. By differentiating between positions along the ladder it contributes to cross-sectional wage inequality.[46] Second, it helps to explain sustained earnings losses after displacement.[47]

It is natural to think how the ladder operates and how it affects workers might vary across types. Some recent research on job ladders has explored the role of worker heterogeneity in terms of productivity, and in empirical implementations sometimes gender and education. We next briefly discuss findings on heterogeneity in these dimensions, and then turn to some first evidence on characteristics and implementations of the ladder by a worker's labor market type.

[43] Moscarini & Postel-Vinay (2018) attribute this to time aggregation bias.

[44] Ozkan et al. (2023) rank firms based on job switchers' earnings growth. See also Barlevy (2008).

[45] Audoly et al. (2024) provide evidence from job adds validating this assumption.

[46] Several authors have used structural job ladder models to quantify this role. See Lentz (2024) for a recent overview.

[47] Characteristics of the ladder also affect an economy's allocative efficiency, as well as the speed of employment recovery after a recession.

7.1 Heterogeneous job ladders in worker productivity

Bagger et al. (2014), with the objective of decomposing life cycle wage growth into the contributions of human capital and job search, allow for heterogeneity across education and skill groups. Using Danish administrative matched employer-employee data, they

find that on-the-job search, i.e. a job ladder, matters more for wage growth of less educated workers, particular early in their careers. Their structural estimation also implies higher job finding rates for more skilled workers, both on and off the job.[48] This generates sorting of more able workers to better employers.

Ozkan et al. (2023) explicitly focus on heterogeneity in job ladder dynamics as a source of differences in wage growth and in lifetime earnings. Using US Social Security Administration (SSA) data for a sample with strong labor force attachment, they document that workers who end up with low lifetime earnings switch jobs more frequently, but experience lower earnings growth when making such a switch. The annual nature of SSA data does not allow for a precise measurement of infra-annual non-employment spells, but the large number of switches with earnings losses of more than 25% suggest that many such switches involve transition through a non-employment spell. Ozkan et al. (2023) find similar patterns in monthly SIPP data.

Since the SSA does not allow for directly measuring high-frequency flow rates, they structurally estimate a life cycle model building on Bagger et al. (2014) with heterogeneity in initial human capital and in learning ability. Their estimates imply that job loss risk declines steeply with lifetime earnings in the lower half of the lifetime earnings distribution, while job finding and contact rates increase with lifetime

[48] Bagger & Lentz (2019) show that in a setting with complementarity of worker skill and firm productivity, more productive workers endogenously expend more effort searching for better jobs.

earnings. They also find that separation rates decline steeply with age, while job finding and contact rates increase. Older workers and those with high lifetime earnings nevertheless make fewer job-to-job transitions because fewer offers dominate their current job. As a result of the large differences, Ozkan et al. (2023) attribute more than 70% of lifetime wage growth differences between bottom and median lifetime earnings workers to differences in unemployment and job ladder risk. More frequent unemployment prevents low-lifetime earnings workers from climbing the job ladder and from accumulating human capital, while lower contact rates imply that they climb the ladder more slowly.

Borovičková & Macaluso (2024) ask a similar question, but allow firms to differ not only in productivity, but also in the learning environment they offer.[49] Using administrative data on Austrian men, they find that while low lifetime earnings workers switch employers more frequently, they do not climb a ladder. Instead, each employer tends to pay less than the previous one. Similar to the previous two papers, they also find that low earnings workers lose their jobs more frequently.

All of this work shows significant heterogeneity in labor market flows as a function of worker heterogeneity in terms of productivity and age.[50] Patterns are consistent in that low lifetime-earnings workers switch employers more frequently, but nonetheless climb the ladder more slowly due to their very high separation rate.

We next explore heterogeneity in the job ladder as a function of a broader notion of labor market types, as estimated above.

[49] Arellano-Bover & Saltiel (forthcoming) document such differences using matched employer-employee data from Brazil and Italy.

[50] Kaas et al. (2023) further consider the role of wealth in a setting with incomplete markets.

7.2 Heterogeneous job ladders in worker labor market type

To our knowledge, GMW and HK are the only papers that analyze how career progression differs with workers' labor market type, broadly defined. Having estimated workers' types as described above, GMW use a structural model with learning about match quality and on-the-job search building on Menzio & Shi (2011) and Menzio et al. (2016) to analyze why flows differ by type. Their quantitative analysis allows the different labor market types to differ in a broad range of model parameters. Results imply that γs face a match quality distribution with a fat right tail. This encourages them to keep sampling and thus leads to higher transition rates both to unemployment and to new jobs, thus accounting for their much shorter job duration. The main inferred reason why γs find jobs more slowly consists in their lower productivity relative to the value of unemployment.[51]

HK's approach captures a very short ladder via the two distinct latent employment types they estimate: short-term and long-term employment.[52] Long-term jobs are characterized by a lower separation rate. There is a ladder in the sense that estimates imply that flows from short-term to long-term employment are fairly common, whereas the reverse transition is rare. The ergodic distributions of labor market types across these states differ. Surprisingly, High-N and High-U types, who spend a lot of time out of work, overwhelmingly work in long-term jobs when working, with ratios of long-term to short-term employment of four and higher. For High-E types, in contrast, this ratio is only 2.7 for women and 1.4 for men. The model specifies that

[51] Gendron-Carrier (2024) estimates a dynamic Roy model of career choice with unobserved heterogeneity (fixed number of labor market types), using Canadian CEEDD data. The analysis concerns the determinants of entry into entrepreneurship, in particular the role played by heterogeneous job ladders during early career work. He finds significant differences in job ladders across types, affecting both entry and the type of entrepreneurial activity.

[52] CPS data, like those in similar short panels, do not contain firm identifiers and therefore do not allow measuring firm characteristics.

All-E types only work in long-term jobs. In terms of flows, High-N and High-U types are much more likely to advance from short-term to long-term jobs than to make the reverse transition. The gap between these flow rates is smaller for high-E types. We conjecture that these surprising differences across types might be a consequence of the short panel, for the same reason as the fairly high E-NE transition rates of the High-E types discussed in Section 3.2.3.

We find very similar patterns in the C-LFS and, for reasons of space, do not present them here. Instead, we focus on patterns in the BEAM. Our approach is to take a mostly descriptive route and to directly document various aspects and forms of a job ladder by labor market type, in order to inform future theories.

To do so, we extend our basic modeling framework of labor force transitions of Section 2 to accommodate job ladders, by simply expanding the number of employment states to differentiate employment by rung on the ladder. Since the BEAM links workers to firms, we are able to obtain firm-level worker counts and average earnings, so that we can consider ladders in terms of both mean firm wage and size. In order to obtain reliable indicators of firm size and wages, we use the full annual sample of workers and firms, instead of the monthly 1% random sample described in Section 3.1.2, and which underlies most of our analysis.

To obtain a firm's rung on the wage ladder while accounting for worker composition, we residualize earnings using time effects, gender, and a cubic age profile and average the residuals within firms. As is typical for administrative data, our "wages" are in fact annual labor earnings. They thus also capture variation in time worked. With this understanding, we still refer to "wages" and to a "wage ladder" in the remainder of this section.

For the size ladder, we consider five categories (< 10, 11-50, 51-1000, 1001-10000,

	Firm Size Category					
Firm Wage Quintile	1	2	3	4	5	Total
1	2.98	4.62	6.25	2.76	3.39	20
2	2.49	3.82	6.57	3.95	3.18	20
3	1.93	3.03	5.98	5.85	2.84	20
4	1.14	1.73	4.16	5.31	8.02	20
5	1.56	1.79	5.48	6.23	4.94	20
Total	10.10	15.00	28.43	24.10	22.37	100

The firm size categories are (<10, 11-50, 51-1,000, 1,001-10,000, >10,000).

Table 11: Joint distribution across job ladder, both genders (in %)

>10000) for the time average of firm-level worker counts.

To compute job distributions by earnings quintiles (our definition of the wage ladder rungs) and size categories, we revert to the 1% random sample. Using this sample, we compute average monthly transitions not only between employment and non-employment, but also between rungs of the ladder. We do so both in the aggregate and separately by labor market type. labor market dynamics are then described by a type-specific Markov chain based on these transitions.[53,54] Table 11 displays the joint distribution of workers over rungs of the wage and size ladders.

The vast majority of jobs in Canada are in firms with over 50 employees. Larger firms tend to pay higher wages, although the association is far from perfect. In what follows, we focus on a one-dimensional firm ladder at a time, either according to size or wage.

[53]Time aggregation is an important issue in this literature, since some sources, in particular the LEHD, do not contain monthly information. Note also that we treat job-to-job flows intermediated by a non-employment spell lasting less than a month as poaching.

[54]We keep the types assigned to individuals based on the moments described in Section 3.1.2. For an analysis of the ladder, an alternative in the spirit of Bonhomme et al. (2022) would be to measure types using moments that contain information about the ladder.

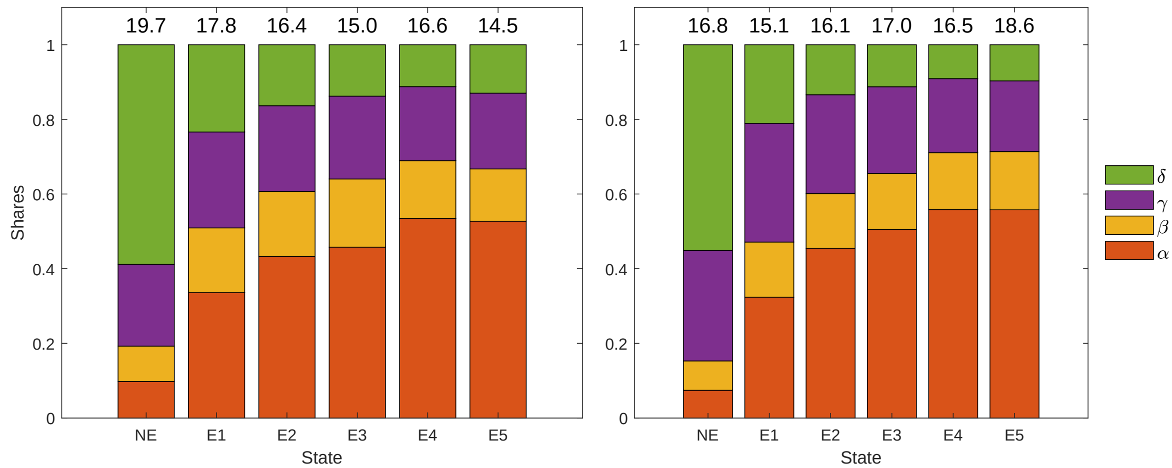

(a) Women (b) Men

Figure 10: Distribution across employment states (top figures, in %), and type composition along the wage ladder

7.2.1 A static snapshot

Next, we document how the distribution of employment over rungs of the ladder varies by type and gender. Figure 10 displays this distribution for the wage ladder, separately for men and women. Each bar shows the type distribution of individuals in that state (a rung of the ladder or non-employment). The numbers at the top of each bar show the share of men or women in that state.

[Figure 10.1 and 10.2 here: Horizontally in a row. Below 10.1 (a) Women Below 10.2 (b) Men Caption: Distribution across employment states (top figures, in %), and type composition along the wage ladder]

Because we have defined wage quintiles using the population, the share of employment conditional on gender need not equal one-fifth on each rung of the wage ladder. Indeed, male employment monotonically increases across wage rungs, whereas they decrease among women except for the fourth wage quintile.

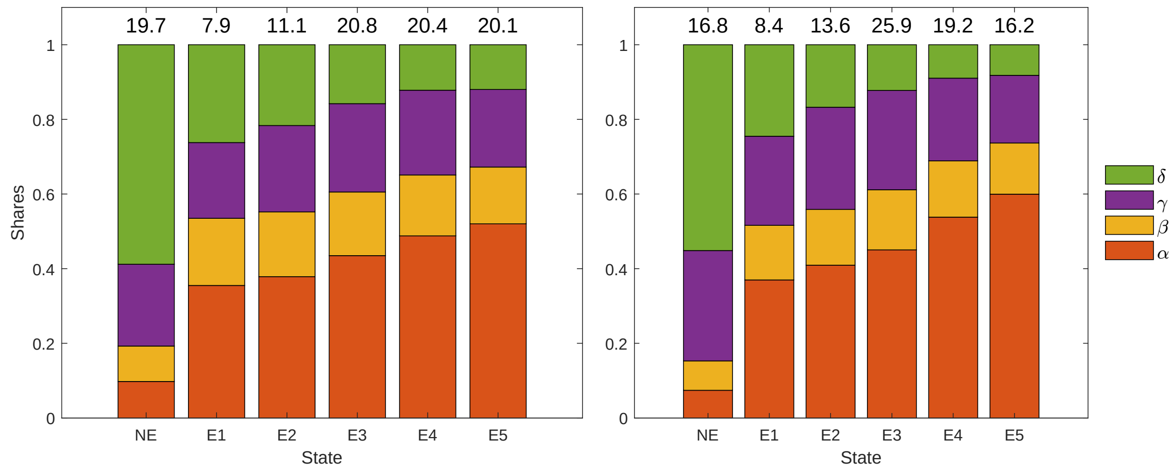

(a) Women

(b) Men

Figure 11: Distribution across employment states (top figures, in %), and type composition along the size ladder

The figure also shows that the distribution of employment over rungs of the ladder differs by type. The share of the stable type α increases monotonically with the wage rung, reaching about one half in the top two rungs. In contrast, γs and particularly δs are not only over-represented among the non-employed, as already shown in Table 5 above, but also on lower rungs of the ladder. As shown in Figure 11, the more stable types are also more common on the top rungs of the size ladder.[55]

[Figure 11.1 and 11.2 here: Horizontally in a row. Below 11.1 (a) Women Below 11.2 (b) Men Caption: Distribution across employment states (top figures, in %), and type composition along the size ladder]

[55]The main feature that sets apart the size distribution from the wage distribution is that the employment shares are larger in the higher rungs. This follows from our definition of the size rungs. One noticeable gender difference is that women tend to be disproportionally employed in larger firms, while we do not see large gender differences across the ladder defined by wage differences.

7.2.2 The job ladder by type

To assess whether the evidence is supportive of job ladders, we rely on Haltiwanger et al.'s (2018) notion of the net poaching rate conditional on job-to-job moves. They build on the following identity describing the net job flows towards firm rung $i \in \{1, \ldots, 5\}$ between any two periods, which we consider for each individual worker type ω:

$$NJF_i(\omega) = H_i^p(\omega) - S_i^p(\omega) + H_i^0(\omega) - S_i^0(\omega),$$

where $H_i^p(\omega)$ and $S_i^p(\omega)$ are poaching hires from and separations to other firms, respectively, and $H_i^0(\omega)$ and $S_i^0(\omega)$ are the analogous flows involving non-employment. We compute these objects using our state-level (rung-level) data.[56] In terms of our basic framework of Section 2,

$$H_i^p(\omega) = \sum_{j \geq 1, j \neq i} n_j^*(\omega) \Pi_{ji}(\omega) \tag{1}$$

$$S_i^p(\omega) = \sum_{j \geq 1, j \neq i} n_i^*(\omega) \Pi_{ij}(\omega). \tag{2}$$

As in Haltiwanger et al. (2018), our focus is on the net job flow rate of rung i due to poaching,

$$NJFR_i^p(\omega) \equiv \frac{H_i^p(\omega) - S_i^p(\omega)}{n_i^*(\omega)}. \tag{3}$$

The key prediction of job ladder models is that $NJFR_i^p(\omega)$ increases in i. That is, the higher the ladder rung, the more firms' poaching from firms on other rungs

[56] This approach does not capture gross poaching flows within a rung. Haltiwanger et al. (2018) compute these objects directly at the firm level.

exceeds their poaching losses. Figure 12 shows that this is indeed the case, providing clear evidence in support of a wage ladder. This is clearly the case in the aggregate (single type case, dotted black line), for both men and women. The evidence is also consistent with a wage ladder conditional on labor market type. The ladders are, however, very heterogeneous, and appear much steeper for γs and δs. We will return to this point.[57]

Note that a stationary employment distribution ($NJF_i(\omega) = 0$) implies $NJFR_i^p(\omega) = S_i^0(\omega) - H_i^0(\omega)$, i.e. net poaching inflows into rung i must be balanced by net outflows towards non-employment. The fact that $NFJR_1^p(\omega)$ is negative for all types then implies that the lowest rung of the ladder has positive net hires from non-employment, but loses to poaching by higher-rung firms. This is particularly pronounced for the unstable worker types.

The bars in Figure 12 decompose the net inflow rates into the poaching inflow and outflow rates, $H_i^p(\omega)/n_i^*(\omega)$ and $S_i^p(\omega)/n_i^*(\omega)$. These flows are small for αs, especially at the top of the ladder. In contrast, the γs are extremely mobile, with significant flows both in and out of just about every employment rung. Gross flow rates are smaller for women, in particular for γs.

[Figure 12.1 and 12.2 here. Horizontally in a row Below 12.1 (a) Women Below 12.2 (b) Men Caption: Poaching Rates along the wage ladder.]

As we have mentioned, the theoretical literature on job ladders generates transition matrices Π that imply $NJFR_i^p$ increasing in i. But note that when using group-level data, $NJFR_i^p$ can also increase in i for mechanical reasons if $n_i^*(\omega)$ decreases in i, even if Π does not imply a ladder.[58] This issue does not arise when rungs

[57] To be concise, we do not show gross poaching rates. These are higher for stable types, who experience non-employment less frequently, and increase with ladder rung except for the top rung.

[58] Consider an extreme case with two rungs, representing 99% and 1% of employment respectively,

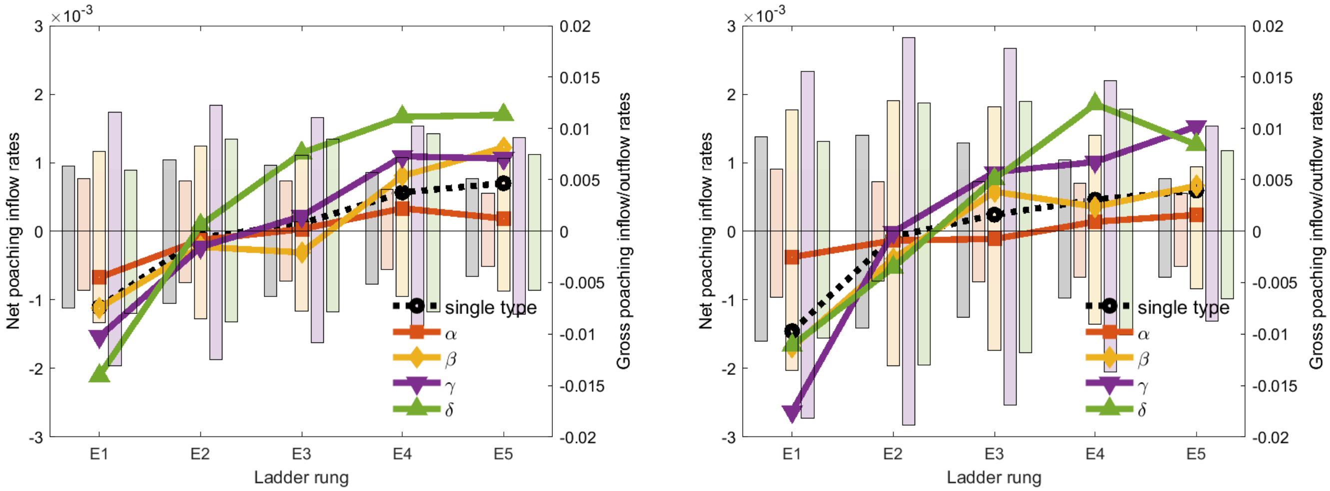

(a) Women (b) Men

Figure 12: Poaching rates along the wage ladder

of the ladder have equal size, e.g. for our single type ladder with rungs defined by wage quintiles, but needs to be reckoned with when heterogeneity implies that $n_i^*(\omega)$ varies with ω. Employment by rung also typically is not constant for size ladders, since size categories are not normally chosen to be equal-sized.

To abstract from this effect and focus on the ladder as implied by the poaching transitions in Π, we compute a constant-employment net poaching rate, setting $n_i^*(\omega)$ in equation (3) equal to the average employment share of each rung (20%) for all i and ω. This is shown in Figure 13. This adjustment accentuates the increase in the net poaching rate with rung i when considering all men jointly (single type, dotted black line), especially towards the lower rungs. The common n_i^* counteracts the attenuating effect of greater employment shares of men in higher rungs (see

and a symmetric transition matrix between these rungs with off-diagonal element p. ($n_1^* > n_2^*$ can still arise in this case if rung 1 has a greater inflow rate from non-employment.) Then $NJFR_2^p = (.99p - .01p)/.01 > NJFR_1^p = (.01p - .99p)/.99$.

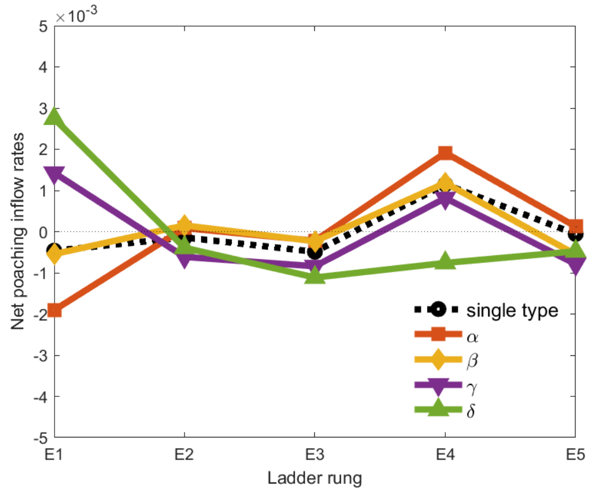

(a) Women

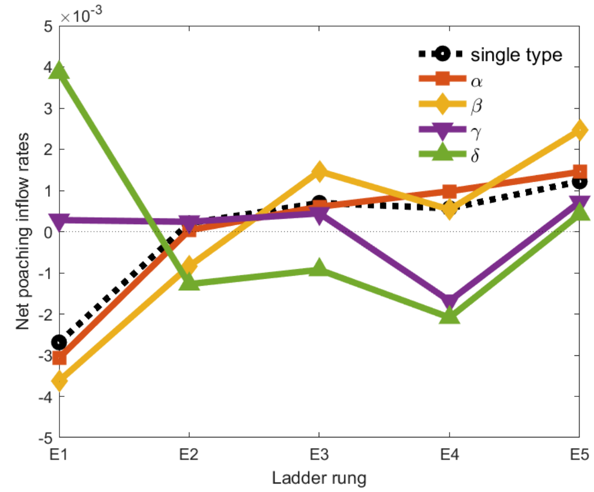

(b) Men

Figure 13: Constant-employment net poaching rates along the wage ladder

Figure 10b) on the slope of $NJFR_i$. For women, the opposite occurs, and constant-employment net poaching hardly varies across rungs.

[Figure 13.1 and 13.2 here. Horizontally in a row Below 13.1 (a) Women Below 13.2 (b) Men Caption: Constant-employment net poaching rates along the wage ladder]

This approach also helps us understand the heterogeneity in the wage ladder across labor market types. It reveals that the greater slope of $NJFR_i(\omega)$ for γs and δs observable in Figure 12 is entirely due to their greater employment shares at the bottom of the wage ladder. Figure 13 reveals that only stable types systematically move up the wage ladder. Figure 14 displays evidence on the size ladder. The single type case is consistent with larger firms being more attractive, however the evidence is weaker compared to wages.

It is also not consistent across labor market types, with evidence for a size ladder

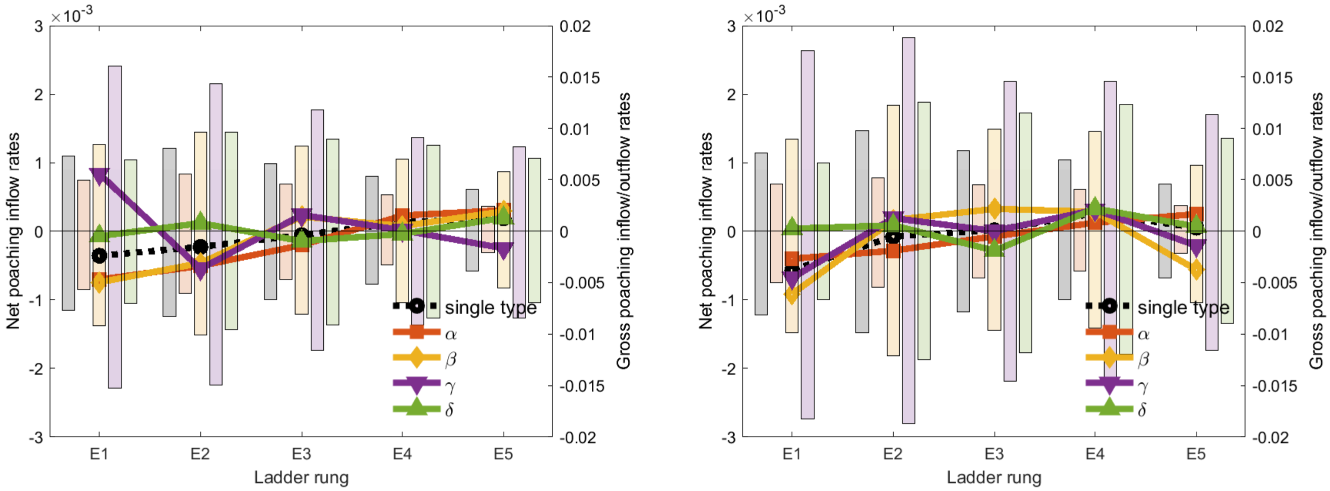

(a) Women

(b) Men

Figure 14: Poaching rates along the size ladder

present only for αs and female βs. In terms of gross flows, the size ladder is overall similar to the wage ladder. The main differences consist in greater gross flows of women of type γ into and out of small firms, compared to low-wage firms.

Figure 15 shows constant-employment net poaching rates for the size ladder. Here, large net poaching rates of medium-sized firms (between 51 and 1,000 employees) stand out.[59] Apart from this, there also is evidence for a size ladder for αs and female βs, plus a weaker pattern for γs and male βs.

To summarize, we find some evidence of both a wage and a size ladder in the aggregate. Dissecting the data by types reveals that this is driven by stable types.

[Figure 14.1 and 14.2 here. Horizontally in a row Below 14.1 (a) Women Below 14.2 (b) Men Caption: Poaching rates along the size ladder]]

[59] While this group accounts for a large share of employment (see Figure 11), this does not affect the constant-employment net poaching rates shown here.

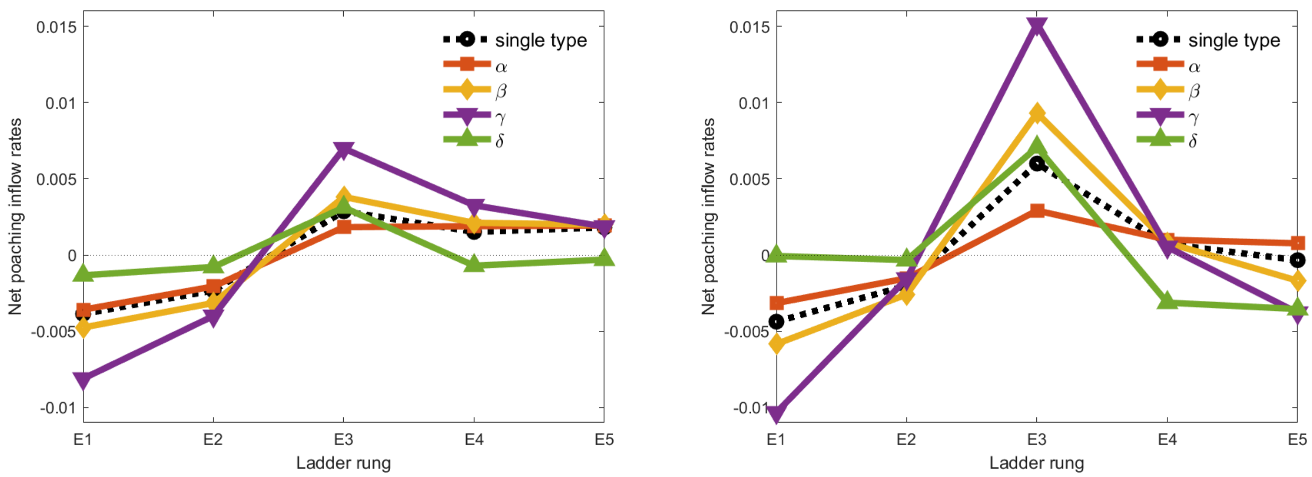

(a) Women

(b) Men

Figure 15: Constant-employment net poaching rates along the size ladder

[Figure 15.1 and 15.2 here. Horizontally in a row Below 15.1 (a) Women Below 15.2 (b) Men Caption: Constant-employment net poaching rates along the size ladder]

7.2.3 Importance of the job ladder versus non-employment flows

Why are αs more prevalent on higher rungs of the ladder, and γs and δs on lower rungs, as shown in Figure 10? How the distribution of employment over rungs of the ladder differs across types depends not only on how each type climbs the ladder, but also on the rungs at which they enter the ladder, as well as how often they fall off the ladder, and from which rungs.

Figure 16 shows how these factors contribute to differences in the employment distribution between αs and γs over rungs of the wage ladder. To that end, the figure shows both the actual distribution (red squares for αs, green triangles for γs) as well as three counterfactuals.

The first counterfactual illustrates how the initial placement on the job ladder affects the stationary distribution across its rungs. For this, we replace the job finding rates across the different rungs of the ladder for the γs with those observed from the αs. Formally, we replace $\Pi_{0i}(\gamma)$ with $\Pi_{0i}(\alpha)$, for each rung $i \in \{1, \ldots, 5\}$ of the wage ladder. These probabilities are then normalized so that the overall job-finding probability of the γs remains unchanged.[60] The figure shows the stationary distribution over rungs of the ladder induced by this counterfactual transition matrix as a dashed green line.[61]

The second counterfactual considers how important differences in separation rates across types are. For this, we replace $\Pi_{i0}(\gamma)$ with $\Pi_{i0}(\alpha)$ for $i \geq 1$, so that the γs face the same separation rates to non-employment as the αs. The remaining probabilities are adjusted so that the γs' rung-to-rung transition probabilities conditional on not separating remain unchanged. The ensuing stationary distribution is shown as a dotted line.

The third exercise lets the γs face the job ladder dynamics of the αs. To do so, we replace the upper and the lower triangular portions of the job-to-job transition submatrix of the γs with the values of the αs. Formally, we replace $\Pi_{ij}(\gamma)$ by $\Pi_{ij}(\alpha)$ for $i \geq 1, i \neq j$. We rescale the other two probabilities on each rung, $\Pi_{i0}(\gamma)$ and $\Pi_{ii}(\gamma)$, so that their ratio remains unchanged.

The implied stationary distribution is shown as a short-long-dashed green line.

[60] In this and the following counterfactuals, a normalization of the remaining probabilities is necessary to ensure that the counterfactual transitions satisfy $\sum_j \Pi_{ij} = 1$ for every i.

[61] Ergodic distributions are informative about empirical ones in this setting, as the ergodic distribution induced by $\Pi(\omega)$ matches the empirical employment distribution over rungs of the ladder extremely well.

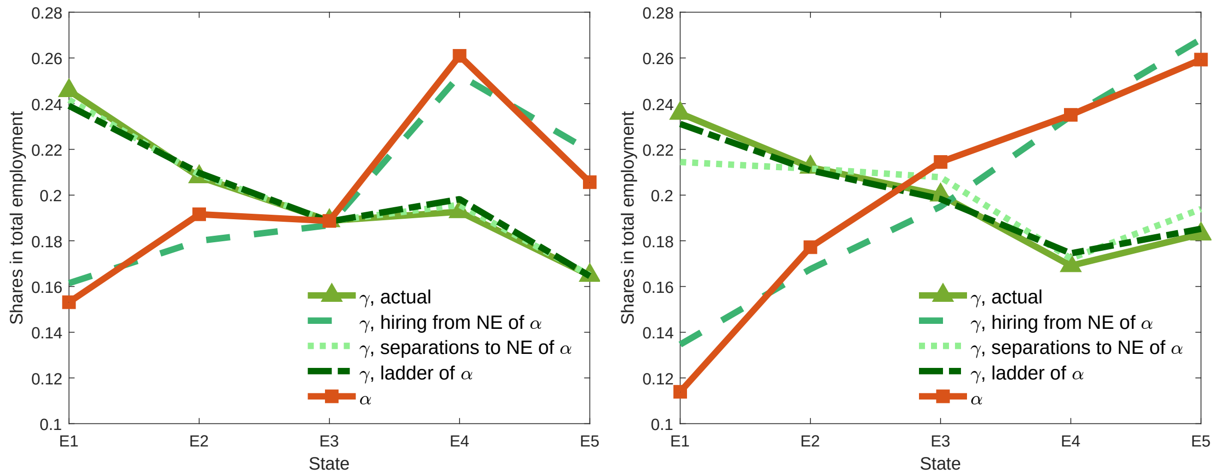

(a) Women (b) Men

Figure 16: Counterfactual employment distributions along the wage ladder

[Figure 16.1 and 16.2 here. Horizontally in a row Below 16.1 (a) Women Below 16.2 (b) Men Caption: Counterfactual employment distributions along the wage ladder]

It is clear from Figure 16 that the first counterfactual, which endows γs with the types of jobs that αs find when leaving non-employment, accounts for almost the entire difference in the distribution across rungs of the ladder between the two types. These differences thus are almost entirely due to the first job after a non-employment spell. Only when the γs face the hiring rates of the αs does their distribution along the ladder increase monotonically, following the same pattern as for the αs. Job ladder dynamics hardly affect differences in the employment distribution across types.[62]

Figure 17 shows the fundamental feature driving this difference, the job finding probabilities from non-employment by rung of the ladder for the different types. For comparison, Figure 18 shows the separation rates to non-employment by rung of the

[62] We repeated the exercise for the size ladder and found similar results, although less pronounced.

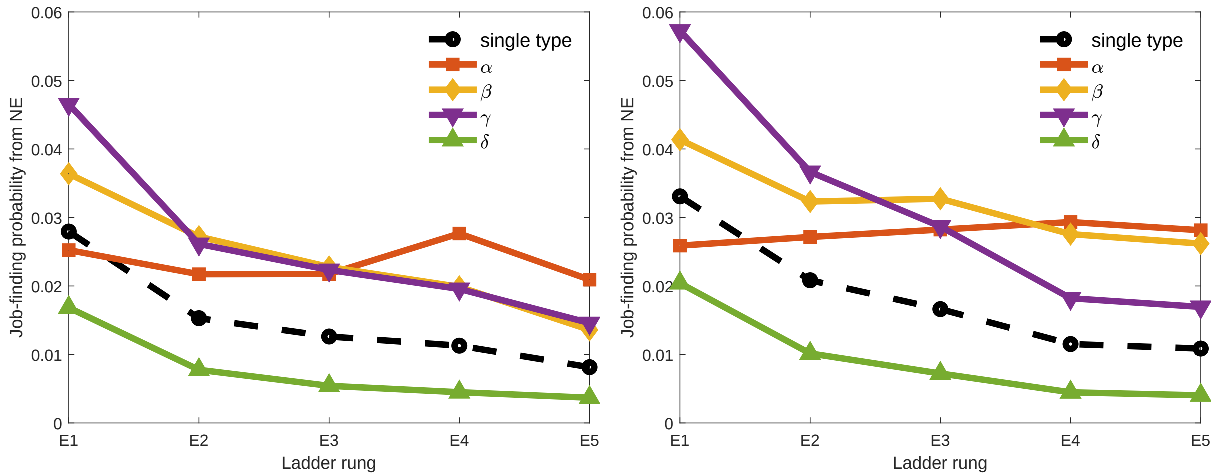

(a) Women (b) Men

Figure 17: Job-finding probabilities from NE along the wage ladder

ladder.

Stable types face a close to uniform rung distribution of their first job out of non-employment. As a consequence, they have a much better starting position on the ladder than unstable types (γs and δs), who tend to enter employment through low-wage jobs, which explains why they are more prevalent in the lower rungs. Separation rates into non-employment mostly differ across types in their level. (Of course, this explains the high non-employment rates of the unstable types, as discussed in Section 3.2.3.) Relative separation rates across rungs of the ladder differ little across types.

To summarize, differences in the distribution of labor market types across the rungs of the job ladder are almost entirely due to where they enter the ladder after a non-employment spell. Differences in the ladder itself contribute little.

[Figure 17.1 and 17.2 here. Horizontally in a row Below 17.1 (a) Women Below 17.2 (b) Men Caption: Job-Finding probabilities from NE along the wage ladder]

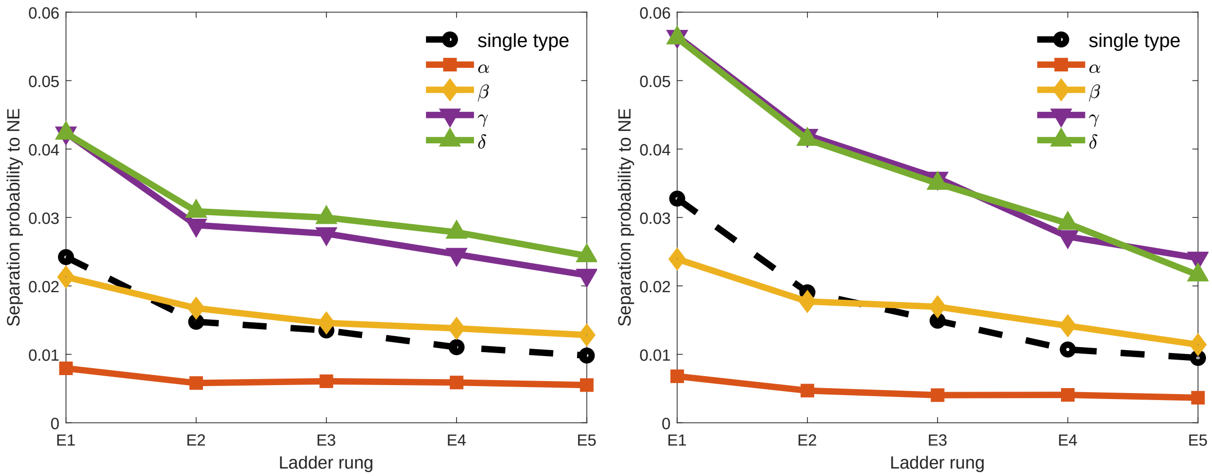

(a) Women (b) Men

Figure 18: Separation probabilities to NE along the wage ladder

[Figure 18.1 and 18.2 here. Horizontally in a row Below 18.1 (a) Women Below 18.2 (b) Men Caption: Separation probabilities to NE along the wage ladder]

7.3 The costs of job loss

We now turn to the consequences of job loss when individuals face a wage job ladder. When wage differentials along the ladder are large, and when it takes time to climb the ladder, wage losses from a layoff can be quite significant, much more so than in models that abstract from ladders and therefore emphasize only the reemployment probability.

The starting point for this literature is the large and persistent earnings losses following layoff from a high-tenure job, as documented initially by Jacobson et al. (1993), Couch & Placzek (2010) or Davis & von Wachter (2011). These initially posed a puzzle to theories of job search: at least in studies of the United States, high job finding rates imply that earnings should recover quickly (see e.g. Low et al. 2010, Davis & von Wachter 2011). Job ladder models have had more success in explaining persistent earnings losses, in particular if separation rates are higher at the bottom of the ladder or after a non-employment spell (Pries 2004, Krolikowski 2017, Jarosch 2023) and jobs at the top of the ladder are very stable (Jung & Kuhn 2019). Indeed, Schmieder et al. (2023) show, using German administrative data, that a large part of earnings losses is due to the fact that displaced workers find new employment in lower-wage firms, in line with job ladder models. Jarosch (2023) finds that wage losses account for more than half of the decline in earnings, and also shows an increase in the risk of future job loss after an initial displacement. Lachowska et al. (2020), in contrast, find a smaller role for firm effects in matched employer-employee panel data from the US state of Washington, and a larger role for reduced match quality.

Recent work has investigated heterogeneity in earnings losses from displacement. An early contribution using survey data from the US Panel Study of Income Dynamics (PSID) is Stevens (1997), who finds that persistent costs arise if displacement is repeated, but that wage losses dissipate for workers who do not experience additional displacement. Flaaen et al. (2019) show that these losses vary by reason for the layoff. Jarosch (2023) finds that the costs of job loss are very heterogeneous. Earnings of the median displaced worker recover after 15 years, while for others, costs are larger and more persistent. Most studies in this literature focus on high-tenure, high-wage workers. Rose & Shem-Tov (2023) change perspective and find that for low-wage earners, the costs of job loss are driven mostly by persistently lower employment and not lower earnings.

GMW document earnings losses from jobs with three or more years of tenure by labor market type in the LEHD data. While losses relative to pre-displacement

State	NE	E1	E2	E3	E4	E5
Wage relative to E5	0	0.150	0.338	0.492	0.656	1

Table 12: Average wages along the wage ladder

earnings are persistent for all types, αs recover most, and most quickly. γs maintain large losses even after 5 years, the maximum horizon considered by GMW. Their structural model attributes this difference across types to the fact that it takes γs very long to find a job that is worth keeping for more than three years. (Recall that only 12% of employment spells of γ workers last more than two years; see Table 6.) The initial job loss will thus be followed by many further transitions.

In the remainder of this section, we provide some illustrative calculations for the cost of displacement in the Canadian BEAM data, by type, using the type-specific transition matrices $\Pi(\omega)$. As in GMW's analysis, the differences in labor market flows and ladders by type have implications for earnings losses.

Table 12 displays the average (residual) earnings within each wage quintile, relative to the wages at the top rung. Wages at the bottom of the ladder are only 15% of those at the top. This large gap overstates wage differences both because it also reflects time worked, and also because firms in which workers work systematically fewer hours will be overrepresented at the bottom of the ladder. In addition, unobserved individual-level productivity differences across firms are imperfectly controlled for. Nevertheless, the wage differences in Table 12 are comparable to those reported by Song et al. (2018) from US Social Security data.[63] The numbers strongly suggest that climbing the wage ladder is associated with substantial earnings gains.

[63] From their Table I, the 10/90 percentile ratio in (non-residualized) average employment-weighted annual earnings across firms ranges from 0.3 in 1981 to 0.2 in 2013, for a sample that already excludes individuals with very low annual earnings.

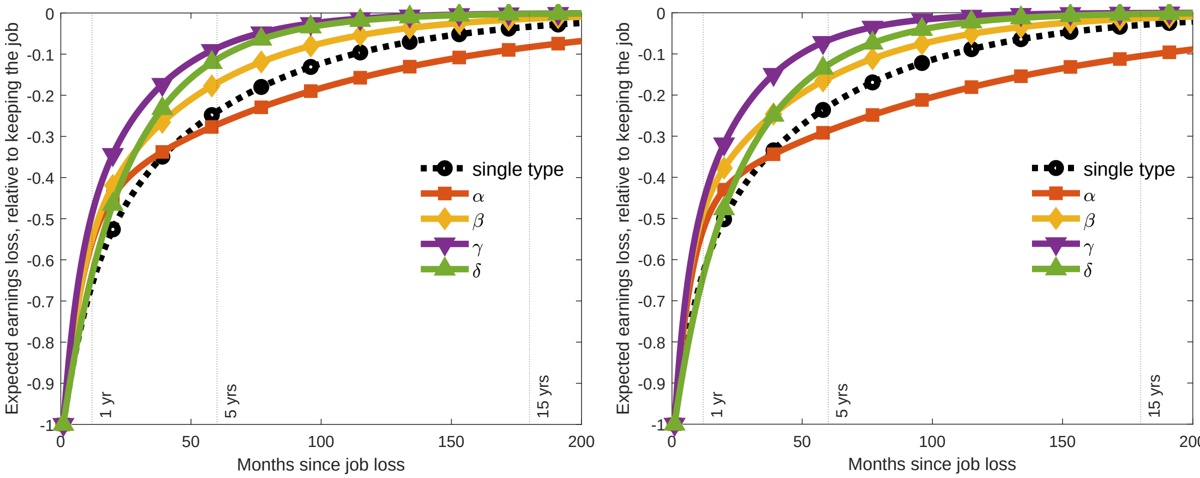

(a) Women (b) Men

Figure 19: Expected earnings loss following layoff from top wage quintile job

Based upon these numbers, we compute the time path of the expected wage loss of somebody who is laid off from a top-rung job, relative to the expected wage of somebody who continues in that job in the month of the layoff.[64] These dynamics follow from the relative wages and the type-specific transition matrix $\Pi(\omega)$.

[Figure 19.1 and 19.2 here. Horizontally in a row Below 19.1 (a) Women Below 19.2 (b) Men Caption: Expected earnings loss following layoff from a top wage quintile job]

[Figure 20.1 and 20.2 here. Horizontally in a row Below 20.1 (a) Women Below 20.2 (b) Men Caption: Expected earnings loss following layoff from a top wage quintile job (displaced workers compared to continuer)]

[64] There is debate in the literature about the counterfactual that the earnings of displaced workers should be compared to. Many papers use the earnings of never-displaced workers. This may overstate losses, since workers not displaced in an initial layoff event may still lose employment later on. Krolikowski (2018) shows that the choice of control group materially affects findings. We choose our counterfactual so that our results will reflect the substantial differences in risk of future layoff across types.

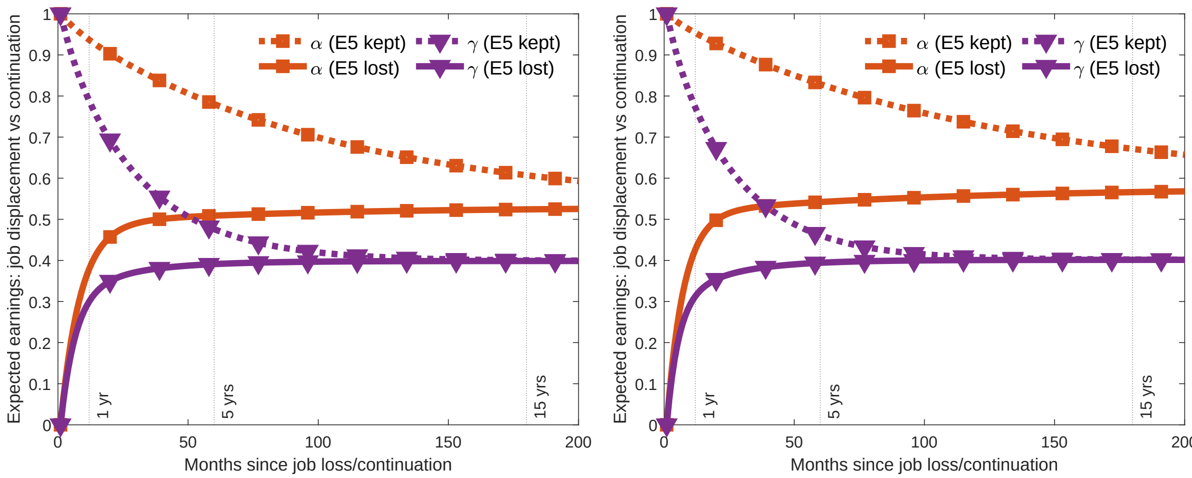

(a) Women (b) Men

Figure 20: Expected earnings following layoff from top wage quintile job (displaced worker compared to continuer)

Figure 19 and Table 13 show that these losses are significant. For the single type, one-year losses amount to 66% for women.[65] That is, one year after the layoff, the expected difference between actual and counterfactual earnings for women is 66% of the top-rung wage. These losses are still significant after five years, at 24%, and only dissipate after 15 years or so, nearly half the length of a typical job market career. Losses for the population average are overall similar. Table 13 also shows that most of the losses, in particular in the long run, are not due to lower employment, but due to lower earnings conditional on employment. That is, not all displaced workers fully climb the ladder again in 15 years.[66] Only for δs, who find jobs more slowly, does a significant share of earnings losses stem from lower employment rates.

[65] Note that by construction, the single type is *not* a simple average of the other types. Table 12 also includes results for the population average.

[66] This analysis only includes post-displacement earnings gains from changing ladder rungs, and excludes career progression of displaced workers while on a rung of the ladder. It may hence overstate earnings losses.

	Women			Men		
Horizon (years after layoff)	1	5	15	1	5	15
Expected earnings loss:						
single type	-0.66	-0.24	-0.03	-0.63	-0.23	-0.03
α	-0.56	-0.27	-0.08	-0.53	-0.29	-0.11
β	-0.55	-0.17	-0.01	-0.49	-0.16	-0.01
γ	-0.49	-0.09	-0.00	-0.46	-0.07	-0.00
δ	-0.63	-0.11	-0.00	-0.64	-0.13	-0.00
total	-0.55	-0.19	-0.04	-0.52	-0.19	-0.05
Expected earnings loss conditional on employment:						
single type	-0.49	-0.22	-0.03	-0.49	-0.21	-0.03
α	-0.43	-0.27	-0.08	-0.43	-0.28	-0.10
β	-0.45	-0.17	-0.01	-0.42	-0.15	-0.01
γ	-0.41	-0.08	-0.00	-0.39	-0.06	-0.00
δ	-0.44	-0.09	-0.00	-0.45	-0.10	-0.00
total	-0.43	-0.18	-0.04	-0.42	-0.19	-0.05

Table 13: Expected earnings losses following a layoff from a top job

Losses are similar across genders, but vary substantially across labor market types. The γs suffer only slightly lower losses than the αs in the short run, but their losses diminish rather quickly over time and are expected to fall below 10% at a five-year horizon, whereas losses for the αs are still 8% even after 15 years.

To better understand these differences, Figure 20 shows the levels of expected earnings relative to the top-rung wage for the displaced worker and the continuer who serves as a comparison, with an emphasis on the αs and the γs. This figure reveals that post-displacement losses of αs quickly stabilize, as they find employment again. But they stabilize at a level significantly below the top wage rung, since αs obtain good starting positions on the ladder (see Figure 17), but then climb the ladder only very slowly. Losses compared to the non-displaced worker nonetheless diminish over time, since the control worker also faced displacement risk.

Both levels and dynamics differ for γs. They also find employment very quickly, but on lower rungs of the ladder. Their loss compared to the continuer nevertheless declines rapidly, since continuing γs also experience high rates of job loss. Convergence would likely be slower for high-tenure workers, since even separation rates for high-tenure γs are much lower than those for the average γ, as shown in Section 5. This is one reason for the difference between the patterns shown here and those documented by GMW. Another reason lies in the much smaller differences in job finding rates between αs and γs in the Canadian BEAM compared to the LEHD used by GMW. These differences warrant further inquiry.

The heterogeneity in losses shown here aligns well with findings from the recent literature. It raises the question to what extent heterogeneity in losses by wage or by reason of layoff might capture heterogeneity in labor market type, since it seems plausible that types differ in their propensity to be engaged in low-wage jobs, or in the reason of layoff. Overall, while precise patterns differ across studies, it is clear that earnings losses due to displacement are heterogeneous. Improving policy responses will require further work to more fully document and understand this heterogeneity.

8 Conclusion

Average job loss and job finding rates, commonly reported in labor market analyses, mask significant heterogeneity. This heterogeneity can be uncovered both in short panels from Labor Force Surveys commonly collected as well as long panels from surveys or administrative sources. A minority of the working-age population accounts for the majority of transitions between employment and non-employment, suggesting that labor market risks and the associated costs are disproportionately concentrated among specific groups of workers.

Recognizing the presence of this heterogeneity will impact many literatures we have not engaged with. This includes inquiries into eg. the costs of business cycles or the effectiveness of active labor market policies. We have explored implications of the heterogeneity uncovered using Canadian data to the literatures on duration dependence in job finding rates, on recall unemployment, and on job ladders. Our analysis suggests that recall employment is a dominant feature in the labor market, particularly among more attached workers. This phenomenon is crucial in understanding the dynamics of job finding rates, as it introduces a different hazard profile compared to non-recall unemployment spells. Moreover, the evidence indicates that dynamic selection and unobserved heterogeneity play a minor role in generating observed duration dependence in non-recall job finding rates, pointing instead to the significance of genuine duration dependence (GDD).

The concept of job ladders also emerged as a critical factor in understanding labor market flows, with heterogeneous patterns evident across different worker types. Low-attachment workers typically start at the bottom of the ladder and climb slowly, if at all, whereas high-attachment workers often access higher rungs directly, maintaining more stable employment over time.

We see several avenues for future research. One key area is the relationship between short and long panel data in measuring labor market transitions, as differences in findings between these two methods suggest that each may capture distinct aspects of labor market dynamics. Additionally, to arrive at a more comprehensive understanding of how workers are impacted by labor market risks requires integrating job ladder models as risk is differentially distributed across rungs of the ladder and types. It will also require engaging with endogenous mobility and consumption smoothing, issues we have abstracted from in our simple back-of-the-envelope calculations to measure how costs from labor market volatility are distributed in the population.

The heterogeneity observed in labor market transitions calls for targeted interventions that can more effectively support vulnerable worker groups and enhance overall labor market efficiency.

References

Abowd, J. M., Kramarz, F. & Margolis, D. N. (1999), 'High Wage Workers and High Wage Firms', *Econometrica* **67**(2), 251–334.
URL: *https://ideas.repec.org/a/ecm/emetrp/v67y1999i2p251-334.html*

Abowd, J. M. & Zellner, A. (1985), 'Estimating gross labor-force flows', *Journal of Business & Economic Statistics* **3**(3), 254–283.

Ahn, H. J. (2023), 'The role of observed and unobserved heterogeneity in the duration of unemployment', *Journal of Applied Econometrics* **38**(1), 3–23.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.2940*

Ahn, H. J. & Hamilton, J. D. (2020), 'Heterogeneity and unemployment dynamics', *Journal of Business & Economic Statistics* **38**(3), 554–569.

Ahn, H. J., Hobijn, B. & Sahin, A. (2023), The Dual U.S. Labor Market Uncovered, NBER Working Papers 31241, National Bureau of Economic Research.
URL: *https://ideas.repec.org/p/nbr/nberwo/31241.html*

Alba-Ramírez, A., Arranz, J. M. & Muñoz-Bullón, F. (2007), 'Exits from unemployment: Recall or new job', *Labour Economics* **14**(5), 788–810.

Albrecht, J., Carrillo-Tudela, C. & Vroman, S. (2018), 'On-the-job search with match-specific amenities', *Economics Letters* **162**, 15–17.

Allman, E. S., Matias, C. & Rhodes, J. A. (2009), 'Identifiability of parameters in latent structure models with many observed variables', *The Annals of Statistics* **37**(6A), 3099–3132.
URL: *http://www.jstor.org/stable/25662188*

Altonji, J. G. & Blank, R. M. (1999), 'Race and gender in the labor market', *Handbook of labor economics* **3**, 3143–3259.

Altonji, J. G., Smith Jr, A. A. & Vidangos, I. (2013), 'Modeling earnings dynamics', *Econometrica* **81**(4), 1395–1454.

Alvarez, F., Borovičková, K. & Shimer, R. (2023), 'Decomposing duration dependence in a stopping time model', *Review of Economic Studies* p. rdad109.

Arellano-Bover, J. & Saltiel, F. (forthcoming), 'Differences in on-the-job learning across firms', *Journal of Labor Economics* .

Audoly, R., Bhuller, M. & Reiremo, T. A. (2024), 'The pay and non-pay content of job ads', *arXiv preprint arXiv:2407.13204* .

Bagger, J., Fontaine, F., Postel-Vinay, F. & Robin, J.-M. (2014), 'Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics', *American Economic Review* **104**(6), 1551–1596.

Bagger, J. & Lentz, R. (2019), 'An empirical model of wage dispersion with sorting', *The Review of Economic Studies* **86**(1), 153–190.

Barlevy, G. (2008), 'Identification of search models using record statistics', *Review of Economic Studies* **75**(1), 29–64.

Bartelsman, E., Haltiwanger, J. & Scarpetta, S. (2013), 'Cross-Country Differences in Productivity: The Role of Allocation and Selection', *American Economic Review* **103**(1), 305–34.
URL: *https://ideas.repec.org/a/aea/aecrev/v103y2013i1p305-34.html*

Bender, R. (2016), Measuring employment and unemployment in canada and the us – a comparison, Technical Report 2015002, Labour Statistics:Technical Papers.

Bonhomme, S., Lamadon, T. & Manresa, E. (2022), 'Discretizing unobserved heterogeneity', *Econometrica* **90**(2), 625–643.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA15238*

Bonhomme, S. & Manresa, E. (2015), 'Grouped patterns of heterogeneity in panel data', *Econometrica* **83**(3), 1147–1184.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA11319*

Borovičková, K. & Macaluso, C. (2024), 'Heterogeneous job ladders', *mimeo, Federal Reserve Bank of Richmond* .

Brown, C. & Medoff, J. (1989), 'The Employer Size-Wage Effect', *Journal of Political Economy* **97**(5), 1027–1059.
URL: *https://ideas.repec.org/a/ucp/jpolec/v97y1989i5p1027-59.html*

Browning, M. & Carro, J. M. (2014), 'Dynamic binary outcome models with maximal heterogeneity', *Journal of Econometrics* **178**(2), 805–823.

Burdett, K. & Mortensen, D. T. (1998), 'Wage differentials, employer size, and unemployment', *International Economic Review* **39**(2), 257–273.
URL: *http://www.jstor.org/stable/2527292*

Cahuc, P., Postel-Vinay, F. & Robin, J.-M. (2006), 'Wage bargaining with on-the-job search: Theory and evidence', *Econometrica* **74**(2), 323–364.

Clark, K. B. & Summers, L. H. (1979), 'Labor market dynamics and unemployment: A reconsideration', *Brookings Papers on Economic Activity* **1979**(1), 13–72.

Couch, K. A. & Placzek, D. W. (2010), 'Earnings losses of displaced workers revisited', *American Economic Review* **100**(1), 572–589.

Davis, S. J. & von Wachter, T. (2011), 'Recessions and the costs of job loss', *Brookings Papers on Economic Activity* **2011**(2), 1–72.

Dey, M. S. & Flinn, C. J. (2005), 'An equilibrium model of health insurance provision and wage determination', *Econometrica* **73**(2), 571–627.

Donovan, K., Lu, W. J. & Schoellman, T. (2023), 'Labor market dynamics and development', *Quarterly Journal of Economics* **138**(4), 2287–2325.

Eliason, M. & Storrie, D. (2009), 'Does job loss shorten life?', *The Journal of Human Resources* **44**, 277 – 302.
URL: *https://api.semanticscholar.org/CorpusID:154922278*

Elsby, M. W., Hobijn, B. & Şahin, A. (2015), 'On the importance of the participation margin for labor market fluctuations', *Journal of Monetary Economics* **72**, 64–82.

Eriksson, S. & Rooth, D.-O. (2014), 'Do employers use unemployment as a sorting criterion when hiring? evidence from a field experiment', *American Economic Review* **104**(3), 1014–39.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.104.3.1014*

Faberman, R. J., Mueller, A. I., Şahin, A. & Topa, G. (2022), 'Job search behavior among the employed and non-employed', *Econometrica* **90**(4), 1743–1779.

Fallick, B. & Ryu, K. (2007), 'The recall and new job search of laid-off workers: a bivariate proportional hazard model with unobserved heterogeneity', *The Review of Economics and Statistics* **89**(2), 313–323.

Farber, H. S., Herbst, C. M., Silverman, D. & Von Wachter, T. (2019), 'Whom do employers want? the role of recent employment and unemployment status and age', *Journal of Labor Economics* **37**(2), 323–349.

Farber, H. S., Silverman, D. & Von Wachter, T. (2016), 'Determinants of callbacks to job applications: An audit study', *American Economic Review* **106**(5), 314–318.

Feldstein, M. S. (1975), 'The importance of temporary layoffs: an empirical analysis', *Brookings Papers on Economic Activity* **1975**(3), 725–745.

Flaaen, A., Shapiro, M. D. & Sorkin, I. (2019), 'Reconsidering the consequences of worker displacements: Firm versus worker perspective', *American Economic Journal: Macroeconomics* **11**(2), 193–227.

Flinn, C. J. & Heckman, J. J. (1983), 'Are unemployment and out of the labor force behaviorally distinct labor force states?', *Journal of labor economics* **1**(1), 28–42.

Forsythe, E., Kahn, L. B., Lange, F. & Wiczer, D. (2022), 'Where have all the workers gone? recalls, retirements, and reallocation in the covid recovery', *Labour Economics* **78**, 102251.

Fujita, S. & Moscarini, G. (2017), 'Recall and unemployment', *American Economic Review* **107**(12), 3875–3916.

Gallant, J., Kroft, K., Lange, F. & Notowidigdo, M. J. (2020), 'Temporary unemployment and labor market dynamics during the covid-19 recession', *Brookings Papers on Economic Activity* .

Gendron-Carrier, N. (2024), Prior work experience and entrepreneurship:the careers of young entrepreneurs. McGill University.

Ghayad, R. (2014), 'Long-term unemployment in the great recession'.

Gregory, V., Menzio, G. & Wiczer, D. G. (2021), The alpha beta gamma of the labor market, Working Paper 28663, National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w28663*

Hall, R. E. (1970), 'Why is the unemployment rate so high at full employment?', *Brookings Papers on Economic Activity* **1970**(3), 369–410.

Hall, R. E. & Kudlyak, M. (2022), Churn and stability: The remarkable heterogeneity of flows among employment, job search, and non-market activities in the us population, Technical report, Federal Reserve Bank of San Francisco.

Hall, R. E. & Milgrom, P. R. (2008), 'The limited influence of unemployment on the wage bargain', *American Economic Review* **98**(4), 1653–1674.

Hall, R. E. & Mueller, A. I. (2018), 'Wage dispersion and search behavior: The importance of nonwage job values', *Journal of Political Economy* **126**(4), 1594–1637.

Haltiwanger, J. C., Hyatt, H. R., Kahn, L. B. & McEntarfer, E. (2018), 'Cyclical Job Ladders by Firm Size and Firm Wage', *American Economic Journal: Macroeconomics* **10**(2), 52–85.
URL: *https://ideas.repec.org/a/aea/aejmac/v10y2018i2p52-85.html*

Heckman, J. J. & Singer, B. (1984*a*), 'Econometric duration analysis', *Journal of econometrics* **24**(1-2), 63–132.

Heckman, J. J. & Singer, B. (1984*b*), 'The identifiability of the proportional hazard model', *The Review of Economic Studies* **51**(2), 231–241.

Heckman, J. & Singer, B. (1984*c*), 'A method for minimizing the impact of distributional assumptions in econometric models for duration data', *Econometrica* pp. 271–320.

Helliwell, J. F. & Huang, H. (2011), 'New measures of the costs of unemployment: Evidence from the subjective well-being of 3.3 million americans', *Macroeconomics: Monetary & Fiscal Policies eJournal* .
URL: *https://api.semanticscholar.org/CorpusID:96427289*

Honoré, B. E. (1993), 'Identification results for duration models with multiple spells', *The Review of Economic Studies* **60**(1), 241–246.

Hornstein, A. (2012), Accounting for unemployment: The long and short of it, Working Paper 12-07, FRB Richmond. Available at SSRN: https://ssrn.com/abstract=2191168 or http://dx.doi.org/10.2139/ssrn.2191168.
URL: *https://ssrn.com/abstract=2191168*

Hornstein, A., Kudlyak, M. & Lange, F. (2014), 'Measuring resource utilization in the labor market', *FRB Economic Quarterly* **100**(1), 1–21.

Jacobson, L. S., LaLonde, R. J. & Sullivan, D. G. (1993), 'Earnings losses of displaced workers', *American Economic Review* pp. 685–709.

Jansson, F. (2002), 'Rehires and unemployment duration in the swedish labour market—new evidence of temporary layoffs', *Labour* **16**(2), 311–345.

Jarosch, G. (2023), 'Searching for job security and the consequences of job loss', *Econometrica* **91**(3), 903–942.

Jarosch, G. & Pilossoph, L. (2019), 'Statistical discrimination and duration dependence in the job finding rate', *The Review of Economic Studies* **86**(4), 1631–1665.

Jones, S. R. G. & Riddell, W. C. (2006), 'Unemployment and nonemployment: heterogeneities in labor market states', *The Review of Economics and Statistics* **88**(2), 314–323.

Jones, S. R. & Riddell, W. C. (1999), 'The measurement of unemployment: An empirical approach', *Econometrica* pp. 147–161.

Jovanovic, B. (1979), 'Job matching and the theory of turnover', *Journal of political economy* **87**(5, Part 1), 972–990.

Jung, P. & Kuhn, M. (2019), 'Earnings losses and labor mobility over the life cycle', *Journal of the European Economic Association* **17**(3), 678–724.

Kaas, L., Lalé, E. & Siassi, N. (2023), 'Job ladder and wealth dynamics in general equilibrium'.

Katz, L. F. (1986), 'Layoffs, recall and the duration of unemployment', *NBER Working Paper* **1825**.

Katz, L. F. & Meyer, B. D. (1990), 'Unemployment insurance, recall expectations, and unemployment outcomes', *The Quarterly Journal of Economics* **105**(4), 973–1002.

Kroft, K., Lange, F. & Notowidigdo, M. J. (2013), 'Duration dependence and labor market conditions: Evidence from a field experiment', *Quarterly Journal of Economics* .

Kroft, K., Lange, F., Notowidigdo, M. J. & Katz, L. F. (2016), 'Long-term unemployment and the great recession: the role of composition, duration dependence, and nonparticipation', *Journal of Labor Economics* **34**(S1), S7–S54.

Kroft, K., Lange, F., Notowidigdo, M. J. & Tudball, M. (2019), 'Long time out: Unemployment and joblessness in canada and the united states', *Journal of Labor Economics* **37**(S2), S355–S397.

Krolikowski, P. (2017), 'Job ladders and earnings of displaced workers', *American Economic Journal: Macroeconomics* **9**(2), 1–31.

Krolikowski, P. (2018), 'Choosing a control group for displaced workers', *ILR Review* **71**(5), 1232–1254.

Krueger, A. B., Cramer, J. & Cho, D. (2014), 'Are the long-term unemployed on the margins of the labor market?', *Brookings Papers on Economic Activity* **2014**(1), 229–299.

Kudlyak, M. & Lange, F. (2014), Measuring heterogeneity in job finding rates among the nonemployed using labor force status histories, Technical Report 8663.

Lachowska, M., Mas, A. & Woodbury, S. A. (2020), 'Sources of displaced workers' long-term earnings losses', *American Economic Review* **110**(10), 3231–3266.

Lancaster, T. (1979), 'Econometric methods for the duration of unemployment', *Econometrica: Journal of the Econometric Society* pp. 939–956.

Lentz, R. (2024), 'Models of wages and mobility in frictional labor markets with random search', *Revue économique* **75**(1), 113–145.

Lentz, R. & Tranaes, T. (2005), 'Job search and savings: Wealth effects and duration dependence', *Journal of labor Economics* **23**(3), 467–489.

Lilien, D. M. (1980), 'The cyclical pattern of temporary layoffs in united states manufacturing', *Review of Economics and Statistics* pp. 24–31.

Lise, J. & Robin, J.-M. (2017), 'The macrodynamics of sorting between workers and firms', *American Economic Review* **107**(4), 1104–1135.

Low, H., Meghir, C. & Pistaferri, L. (2010), 'Wage risk and employment risk over the life cycle', *American Economic Review* **100**(4), 1432–1467.

Lucas, R. E. (1987), *Models of Business Cycles*, Basil Blackwell, Oxford.

Machin, S. & Manning, A. (1999), 'The causes and consequences of longterm unemployment in europe', *Handbook of labor economics* **3**, 3085–3139.

Menzio, G. & Shi, S. (2011), 'Efficient search on the job and the business cycle', *Journal of Political Economy* **119**(3), 468–510.
URL: *https://doi.org/10.1086/660864*

Menzio, G., Telyukova, I. A. & Visschers, L. (2016), 'Directed search over the life cycle', *Review of Economic Dynamics* **19**, 38–62. Special Issue in Honor of Dale Mortensen.
URL: *https://www.sciencedirect.com/science/article/pii/S1094202515000241*

Mincer, J. & Jovanovic, B. (1981), Labor mobility and wages, *in* 'Studies in labor markets', University of Chicago Press, pp. 21–64.

Morchio, I. (2020), 'Work histories and lifetime unemployment', *International Economic Review* **61**(1), 321–350.

Moscarini, G. & Postel-Vinay, F. (2008), 'The timing of labor market expansions: New facts and a new hypothesis', *NBER Macroeconomics Annual* **23**(1), 1–52.

Moscarini, G. & Postel-Vinay, F. (2016), 'Did the job ladder fail after the great recession?', *Journal of Labor Economics* **34**(S1), S55–S93.

Moscarini, G. & Postel-Vinay, F. (2018), 'The Cyclical Job Ladder', *Annual Review of Economics* **10**(1), 165–188.
URL: *https://ideas.repec.org/a/anr/reveco/v10y2018p165-188.html*

Mueller, A. I. & Spinnewijn, J. (2023), The nature of long-term unemployment: Predictability, heterogeneity and selection, Technical report, National Bureau of Economic Research.

Nekoei, A. & Weber, A. (2015), 'Recall expectations and duration dependence', *American Economic Review* **105**(5), 142–146.

Nunley, J. M., Pugh, A., Romero, N. & Seals, R. A. (2017), 'The effects of unemployment and underemployment on employment opportunities: Results from a correspondence audit of the labor market for college graduates', *Ilr review* **70**(3), 642–669.

Nüß, P. (2018), Duration dependence as an unemployment stigma: Evidence from a field experiment in germany, Technical report, Economics Working Paper.

Ozkan, S., Song, J. & Karahan, F. (2023), 'Anatomy of lifetime earnings inequality: Heterogeneity in job-ladder risk versus human capital', *Journal of Political Economy Macroeconomics* **1**(3), 506–550.

Poschke, M. (2023), 'Wage employment, unemployment and self-employment across countries', **16271**.

Postel-Vinay, F. & Robin, J.-M. (2002), 'Equilibrium wage dispersion with worker and employer heterogeneity', *Econometrica* **70**(6), 2295–2350.

Poterba, J. M. & Summers, L. H. (1986), 'Reporting errors and labor market dynamics', *Econometrica: Journal of the Econometric Society* pp. 1319–1338.

Pries, M. J. (2004), 'Persistence of employment fluctuations: A model of recurring job loss', *The Review of Economic Studies* **71**(1), 193–215.

Robertson, M. (1989), 'Temporary layoffs and unemployment in canada', *Industrial Relations: A Journal of Economy and Society* **28**(1), 82–90.

Rogerson, R., Shimer, R. & Wright, R. (2005), 'Search-theoretic models of the labor market: A survey', *Journal of Economic Literature* **43**(4), 959–988.

Rose, E. K. & Shem-Tov, Y. (2023), 'How replaceable is a low-wage job?', *NBER Working Paper* **31447**.

Schmieder, J. F., Von Wachter, T. & Heining, J. (2023), 'The costs of job displacement over the business cycle and its sources: evidence from germany', *American Economic Review* **113**(5), 1208–1254.

Shibata, M. I. (2019), Labor Market Dynamics: A Hidden Markov Approach, IMF Working Papers 2019/282, International Monetary Fund.
URL: *https://ideas.repec.org/p/imf/imfwpa/2019-282.html*

Shimer, R. (2005), 'The cyclical behavior of equilibrium unemployment and vacancies', *American Economic Review* **95**(1), 25–49.

Song, J., Price, D. J., Guvenen, F., Bloom, N. & von Wachter, T. (2018), 'Firming Up Inequality', *The Quarterly Journal of Economics* **134**(1), 1–50.
URL: *https://doi.org/10.1093/qje/qjy025*

Sorkin, I. (2018), 'Ranking firms using revealed preference', *The quarterly journal of economics* **133**(3), 1331–1393.

Stevens, A. H. (1997), 'Persistent effects of job displacement: The importance of multiple job losses', *Journal of Labor Economics* **15**(1, Part 1), 165–188.

Sullivan, D. C. & von Wachter, T. M. (2009), 'Job displacement and mortality: An analysis using administrative data', *Quarterly Journal of Economics* **124**, 1265–1306.
URL: *https://api.semanticscholar.org/CorpusID:30813472*

Taber, C. & Vejlin, R. (2020), 'Estimation of a roy/search/compensating differential model of the labor market', *Econometrica* **88**(3), 1031–1069.

Topel, R. H. & Ward, M. P. (1992), 'Job mobility and the careers of young men', *The Quarterly Journal of Economics* **107**(2), 439–479.

Van den Berg, G. J. (2001), Duration models: specification, identification and multiple durations, *in* 'Handbook of econometrics', Vol. 5, Elsevier, pp. 3381–3460.

Van den Berg, G. J. & Van Ours, J. C. (1996), 'Unemployment dynamics and duration dependence', *Journal of Labor Economics* **14**(1), 100–125.

Wang, J. (2010), 'Consistent selection of the number of clusters via crossvalidation', *Biometrika* **97**(4), 893–904.
URL: *https://EconPapers.repec.org/RePEc:oup:biomet:v:97:y:2010:i:4:p:893-904*

Winkelmann, L. & Winkelmann, R. (1998), 'Why are the unemployed so unhappy? evidence from panel data', *Economica* **65**, 1–15.
URL: *https://api.semanticscholar.org/CorpusID:16540146*

Wright, R., Kircher, P., Julien, B. & Guerrieri, V. (2021), 'Directed search and competitive search equilibrium: A guided tour', *Journal of Economic Literature* **59**(1), 90–148.