Latent Worker Dynamics and Wage Changes: A Pocket of Vulnerability*

Hie Joo Ahn†
Federal Reserve Board

Cynthia Doniger‡
Federal Reserve Board

March 27, 2026

Abstract

This paper investigates the effects of latent labor-market heterogeneity on wage changes among new hires. To this end, we uncover latent labor market states and unobserved worker types by estimating a *multi-channel latent Markov model*, using individual histories of labor force status and job starts under various circumstances from the Survey of Income and Program Participation. We identify latent labor-market states that are highly transient, typically experienced by workers with unstable employment who frequently switch jobs or become unemployed. We classify individuals who encounter transient states within a year or less as the vulnerable type, and all others as the stable type. The vulnerable type, comprising 20 percent of the civilian population, does not experience wage gains when switching jobs, with wages showing significant procyclicality at new jobs, in contrast to typical job switchers who experience wage gains and acyclical wage changes. These observations lead us to consider job-switches originating from an unstable employed state as a "godfather shock." We explore these facts in a search model in which some, but not all, workers are willing to sacrifice wages to be employed.

JEL Classifications:
J3: Wages, Compensation, and Labor Costs
E24 Employment, Unemployment, Wages

Keywords: Wage Cyclically, Tenure and Turnover.

*Disclaimer: PRELIMINARY: DO NOT CIRCULATE The views expressed in this paper are those of the authors and do not necessarily reflect the views and policies of the Board of Governors or the Federal Reserve System.

†Federal Reserve Board of Governors, 20th Street and Constitution Avenue NW, Washington, DC 20551, U.S.A. Email: econ.hjahn@gmail.com

‡Email: Cynthia.Doniger@frb.gov

1 Introduction

Heterogeneity has been at the center of discussions in research on labor market dynamics. A growing body of literature has employed sophisticated statistical modeling and machine learning algorithms to uncover latent heterogeneity crucial for explaining these dynamics. Examples are the duration dependence in unemployment hazards (e.g., Ahn and Hamilton, 2020; Mueller and Spinnewijn, 2023), the slow recovery of long-term unemployment after an economic recession (e.g., Shibata, 2019; Gregory, Menzio and Wiczer, 2025); the comprehensive characterization of excess employment persistence, duration dependence in unemployment hazards, and heterogeneous labor-force attachment (e.g., Hall and Kudlyak, 2022; Ahn, Hobijn and Sahin, 2023); and movements along the job ladder and the consequent inequality (Ozkan, Song and Karahan, 2023; Borovičková and Macaluso, 2025). Despite its importance, effects of latent labor-market heterogeneity on wages changes of new hires have been relatively under-explored.

This paper uncovers dynamic latent heterogeneity in individual labor market outcomes by jointly analyzing individual histories of labor force status and job switches. Our focus is to identify pockets of vulnerability among workers and examine their wage changes at a new job. Methodologically, we employ a multi-channel latent Markov model (MLMM), which distills distinct latent labor-market states and their transitions from two parallel individual-level sequences: labor force status and job switches, including reported circumstances such as nonemployment spells between jobs. We use individual-level data from the Survey of Income and Program Participation (SIPP), which, with its rich five-year labor market panels, provides an ideal setting for identifying latent states and worker types. With individuals' histories of experiencing transient states that are prone to nonemployment, we classify workers into vulnerable and stable types, and analyze the extent to which these types predict wage changes among new hires and their cyclicality. We interpret these findings using a structural model of on-the-job search with two allocative shocks, one of which—the godfather shock—induces job switches that fail to raise wages. To our knowledge, this paper is the first to employ an MLMM to uncover crucial worker heterogeneity in the joint variation of wage changes and job-to-job moves, highlighting the role of the godfather shock.

Employer-to-employer (EE, henceforth) transitions capture important heterogeneity within employment that conventional labor force classifications—employment, unemployment, and nonparticipation—are limited to capture. Of particular interest is heterogeneity within EE

transitions, including those that improve a worker's value or wages and those that do not, as well as those that lead to stable employment and those that do not. However, incorporating individual-level histories of EE transitions alongside labor force status into a latent Markov model presents methodological challenges. First, it substantially increases computational costs by expanding the dimension of the latent state sequences to be considered. We tackle the resulting computational and numerical difficulties by combining the Baum–Welch algorithm with an EM algorithm tailored to MLMM. Second, a key identifying assumption of the MLMM—the conditional independence—suggests that, conditional on an individual's latent state, job-switch and labor-force status information should not directly predict past or future outcomes. We address this issue by constructing the job-switch data in a manner that satisfies the identifying assumption, while also allowing for errors in reported histories and inconsistencies between job switches and labor force status. Finally, to account for the non-random nature of frequently observed non-responses in SIPP, we treat missing status as a separate observable category, following the approach of Ahn and Hamilton (2022).

Our MLMM uncovers seven latent labor-market states: two employment-prone (E-prone) states distinguished by the availability of job switch information, a nonparticipation-prone (N-prone) state, a missing-prone (M-prone) state, an unemployment-prone (U-prone) state, and two states characterized by considerable probabilities of being in all three labor force statuses distinguished by recent job switches and holding the same job (Xn and Xs prone, respectively). The U-, Xn-, and Xs-prone states, which together constitute only 8 percent of the population, are highly *transient* relative to other states, exhibiting low own persistence. In particular, workers in the Xn-prone state are employed about half the time and tend to report starting a new job after a long non-employment spell or the end of a previous job. Those in the Xs-prone state are employed roughly a quarter of the time but report either holding the same job or starting a new job after a short unemployment spell. The three transient states are interconnected: a non-negligible share of individuals switch between the three latent states, suggesting that the experience of labor market volatility is concentrated among a small group of workers. Those in these transient states are more likely to have lower educational attainment, shorter tenure at their previous job, and to be overrepresented in agriculture, mining, and retail industries, where seasonal demand is prevalent.

The model, together with the Viterbi algorithm, also uncovers the individual histories of most likely latent states. Leveraging this feature, we classify those who experienced the transient states

within an 11-month interval as the *vulnerable* type during the duration of the interval, while those who do not are classified as the *stable* type. Vulnerable workers tend to have unstable employment or change jobs frequently. We define an 11-month interval—five months before and after the incidence of a transient state—to classify individuals with seasonal jobs or exposure to seasonal risks as part of the vulnerable type. This notion of type identification is akin in spirit to the band-pass filter. Because transient states often occur in succession, the average duration of vulnerability is 16 months, compared with 45 months for the stable type. This approach ensures that the identified type reflects temporal stability, consistent with the notion of a "type," while still allowing for an individual's type to change over time. This feature is distinguished from the previous studies (e.g., Gregory et al., 2025; Borovičková and Macaluso, 2025) based on K-means clustering that do not allow type changes. We externally validate this concept of type using various supplementary wage regressions.

We find that vulnerable workers make up a minority—about 20 percent of the civilian noninstitutional population—but display wage dynamics that differ sharply from those of the stable type and from prior findings in the literature. Gertler, Huckfeldt and Trigari (2020) report that EE transitions generally lead to wage gains relative to prior jobs, and that wages for new hires tend to be acyclical because job changes are largely driven by worker preferences. In contrast, vulnerable workers typically do not experience wage increases and often face wage cuts. Their wages are also significantly procyclical, declining mainly as the unemployment gap widens. Within this group, those with long tenure at their previous job and longer non-employment gaps between two consecutive jobs are more likely to incur wage cuts and to face greater cyclical risk in wages at their new job. It is notable that vulnerable workers with shorter non-employment gaps exhibit significantly procyclical wages, in sharp contrast to the insignificant cyclicality observed among those with longer gaps.

Our results highlight *the pocket of vulnerability* in which workers switch jobs involuntarily, not due to changes in preference, and are therefore more exposed to cyclical and seasonal risks than other job switchers. Moreover, long tenure at the previous job does *not* necessarily protect workers from this vulnerability; rather, for vulnerable workers it tends to have a more negative effect on subsequent wages, likely reflecting both the loss of job-specific human capital and greater reallocation frictions associated with long tenure. In addition, the excess wage cyclicality among vulnerable workers with short non-employment gaps suggests a willingness to accept lower wages in order to stay employed.

We interpret our empirical findings through the lens of a quantitative on-the-job search model with worker and job heterogeneity, incorporating the "godfather shock" in the spirit of Tjaden and Wellschmied (2014). The godfather shock, also referred to as a reallocation shock, is an on-the-job offer that cannot be refused, even if the job switch is unprofitable; refusal results in unemployment. This shock is first introduced by Jolivet, Postel-Vinay and Robin (2006) and incorporated into an on-the-job search model with two-sided heterogeneity by Tjaden and Wellschmied (2014) to study wage dispersion. While a canonical on-the-job search model, as considered by Hornstein, Krusell and Violante (2011), only characterizes EE transitions that improve worker value, the model with godfather shocks can characterize EE transitions that do not improve worker value in terms of wages. Tjaden and Wellschmied (2014) show that incorporating godfather shocks better accounts for the observed incidence and magnitude of wage cuts between adjacent jobs—attributing 60 percent of such cuts to godfather shocks—and significantly reduces the share of wage dispersion attributable to search frictions, thereby supporting a calibration with greater overall wage dispersion. The godfather shock is empirically supported by evidence on on-the-job searches driven by the fear of job loss or accompanied by involuntary separation, as investigated by Fujita (2010), Ahn and Shao (2021), and Ahn and Rudd (2025).

In a cyclical context, Gertler et al. (2020) employ the godfather shock to explain the procyclicality of wages following job-to-job transitions. Their model posits that the share of job-to-job transitions accounted for by godfather shocks rises in recessions and, since such transitions are associated with worse wage outcomes, reduce the observed wage gains for the average job-to-job mover in a recession. Our empirical findings do not support this mechanism. Specifically, we find that the outcomes of workers experiencing godfather shocks are themselves worse while prevalence is only moderately cyclical.

To explain the mechanism behind the empirical findings, we develop a novel structural model that incorporates on-the-job search, godfather shocks, and, importantly, and heterogeneous valuations of leisure among workers coupled with wage posting for the unemployed. The model we develop is similar to Robin (2011), which studies the Bertrand wage setting mechanism in a stochastic environment simplified to include only one firm and one worker type. We modify this framework to include godfather shocks, heterogeneous value of leisure, and wage posting for the unemployed.[1] For clarity of exposition, we posit that workers who face a godfather

[1] Mixtures of these contracts have also been studied by Postel-Vinay and Robin (2004), Doniger (2023b), Doniger

shock continue employment but their wages fall to their reservation wage. In the model, the low-leisure-value workers are employed at higher rates than the high-leisure value counterparts, unsurprisingly. The model generates cyclical reservation wages for the low-value-of-leisure workers while the reservation wages of high-value-of-leisure workers are acyclcial. This creates the wage behavior observed in the data: relatively acyclical wages of new hires from unemployment and new hires following voluntary job-to-job mobility *and* procyclical wages for the typical recipient of a godfather shock.

This paper is methodologically closest to recent studies using hidden Markov models, such as Shibata (2019); Ahn et al. (2023), but differs in several important respects. First, we employ an MLMM that tracks individual histories of both labor force status and job switches—along with their circumstances—whereas previous studies use single-channel models. Second, we focus on the implications of latent states and unobserved worker types for wage changes. To the best of our knowledge, this is the first paper to apply an MLMM to jointly identify latent states and worker types from individual histories of labor force status and job transitions, uncovering stark heterogeneity in the wages and cyclicality of new hires. Finally, we introduce a novel structural model featuring worker and job heterogeneity, on-the-job search, and godfather shocks to interpret the empirical findings. Using the model estimates, we calibrate the share of employer-to-employer (EE) transitions attributable to the godfather shock—a key parameter in job-ladder models—and uncover the link between persistent wage heterogeneity and exposure to such shocks, a contribution that extends the job-ladder literature. A detailed review of related work is provided in Section 2.

The paper is organized as follows. Section 2 provides a literate review. Section 3 describes the data used in the empirical analysis. Section 4 outlines the empirical method and presents the estimation results. Section 5 details the identification of worker types based on individual latent-state histories. Section 6 examines the effects of latent states and worker types on wage changes. Section 7 discusses a structural mechanism that explains the empirical findings. Section 8 concludes.

and Toohey (2022) , and Flinn and Mullins (n.d.).

2 Literature Review

This paper bridges multiple strands of the macro-labor literature, encompassing the role of latent heterogeneity in labor market dynamics, on-the-job search and the job-ladder model, and wage changes.

To begin with, recent studies uncover latent heterogeneity not well captured by observable attributes, using various empirical methods, including statistical models and machine learning. Kudlyak and Lange (2014) examine heterogeneity in the job-finding rates of the nonemployed by classifying them according to their labor force status histories, based on the Current Population Survey. Hall and Kudlyak (2019) identify five distinct types of workers with heterogeneous model parameters, using a search-and-matching framework and standard value function estimation to match the observed data in the CPS. Somewhat differently, Ahn and Hamilton (2020) estimate a nonlinear state space model with an extended Kalman filter to identify the inflows and outflows of workers with unobserved heterogeneity that lead to long-term unemployment using unemployment duration data from the CPS. Mueller and Spinnewijn (2023) employ an ensemble machine-learning model to predict long-term unemployment and assess the role of heterogeneity using Swedish individual-level data. Shibata (2019) estimate a hidden Markov model (HMM) to uncover latent labor market states not well reflected in observed employment, unemployment, and nonparticipation status, based on CPS individual-level data. Ahn et al. (2023) identify a dual labor market structure using an HMM with inequality and zero restrictions that distinguish segmented labor market structures. Gregory et al. (2025) and Borovičková and Macaluso (2025) employ K-means clustering to uncover heterogeneity in individual labor force transitions and EE transitions along the job ladder, respectively. Ozkan et al. (2023) focus on latent heterogeneity in job-ladder dynamics and on-the-job learning as a mechanism of lifetime earnings inequality in the United States.

In particular, the main finding of this paper closely aligns with those of Ahn (2023) and Ahn et al. (2023). Using nonlinear state space models and unemployment duration data disaggregated by reason for unemployment, Ahn (2023) explore dynamic latent heterogeneity within groups of unemployed workers and find that individuals prone to long-term unemployment are present among both permanent job losers and voluntary job leavers, though their share is much larger in the former group. This suggests that some EE transitioners behave similarly to job losers who experienced involuntary permanent separation. Similarly, Ahn et al. (2023) show that

EE transitions are more prevalent in the secondary labor market, characterized by unstable employment, than in the primary market, which represents employment stability.

Next, this paper also speaks to the literature on wage changes and the user cost of labor. A foundational paper, Kudlyak (2014), introduces and empirically estimates the user cost of labor using the National Longitudinal Survey of Youth, demonstrating that it is significantly more procyclical than average or new-hire wages—a finding further highlighted by Kudlyak (2024). Gertler et al. (2020) argue that wages are acyclical once match quality is accounted for. Gertler et al. divide workers into those who transition between jobs quickly and those who take longer, and essentially interpret the quick transitions between employers reveal a preference for the second employer over the first and an imply increase in match quality.[2] Differently, Cairó and Cajner (2018) and Doniger (2023a) explore heterogeneity along the education margin in the wage dynamics. Doniger (2023a) documents heterogeneity separations and in the cyclical sensitivity cost of labor with respect to observables—education and experience—and finds that separations decreases while the cyclical sensitivity of the cost of labor increases as education and experience rise. Doniger argues that the margins of adjustment to shocks vary across groups, with secularly stable groups being less likely to experience unemployment but more likely to experience wage revisions.[3] Ahn and Rudd (2025), based on a Bayesian structural VAR model, identify two reallocation shocks: one associated with job quits that raises wage inflation, and another associated with job losses that lowers wage inflation.

This paper contributes to the literature on search-and-matching models, wage dispersion, and labor market frictions. In the wage inequality literature, the primary focus has been on the relative contributions of worker and job heterogeneity and search frictions to observed wage dispersion, with previous studies reaching divergent conclusions. Postel-Vinay and Robin (2002) show that both productivity differences and search frictions are important. Carrillo-Tudela (2012) find that search frictions explain over 40 percent of wage inequality within skill groups. Hornstein et al. (2011) argue that on-the-job search is the key mechanism behind large frictional wage dispersion, whereas Jolivet et al. (2006) and Tjaden and Wellschmied (2014) emphasize reallocation, or godfather, shocks as the drivers of observed wage cuts and frequent downward

[2]Specifically, the authors mentioned the following: *"In our most narrow measure, an individual is classified as a job changer only if the individual is recorded at a new job with no interruption in earnings. In the second measure, we also classify new hires with a single month of no earnings between jobs as job changers, allowing for the possibility that the worker found the new job from employment but took a short break between job spells. (pg. 1885)"*

[3]Doniger (2023a) applies the adjustment of Fujita, Moscarini and Postel-Vinay (2020) and shows that Cairó and Cajner (2018) holds even when accounting for EE transitions.

moves on the job ladder without requiring large wage-offer dispersion or a dominant role for search frictions, as the shock itself drives much of the lower-side of the wage distribution.[4] While Jolivet et al. (2006) characterize reallocation shocks as leading workers to accept jobs that do not improve value or wages—losses that are front-loaded and quickly offset as workers move up the job ladder—Tjaden and Wellschmied (2014), with a nonstationary (random-walk) idiosyncratic worker shock with persistent effects, demonstrate that such shocks can also lead to value-non-improving moves and job switches that do not result in upward mobility on the job ladder.

The empirical literature on heterogeneity in labor market outcomes also relates to our paper. One strand of such studies focus on differences on worker productivity (e.g., Low, Meghir and Pistaferri, 2010, Doniger and Toohey, 2022). Low et al. (2010) distinguish productivity risk—persistent, influencing wages across jobs—from employment risk, which is transitory and tied to job arrival or destruction, and develop a life-cycle model of consumption, labor supply, and job mobility in the presence of search frictions. Another body highlights the role of job or demand heterogeneity. Abowd, Kramarz and Margolis (1999) find that both worker and job heterogeneity are important, and that job changes across firms explain a significant share of wage variance over time, with firm effects accounting for a substantial part of wage changes when workers switch employers. Tanaka, Warren and Wiczer (2023) empirically show that earnings growth in the transition increases with net employment growth at the destination firm and argue that firm dynamics themselves are key to workers' earnings growth during job-to-job transitions.

Somewhat differently, a body of research focuses on worker heterogeneity. For example, Morchio (2020) demonstrate that most prime-age unemployment spells originate from a small subset of workers, highlighting the importance of persistent, unobserved worker heterogeneity in labor market dynamics; Mueller and Spinnewijn (2023) likewise find that unobserved worker heterogeneity is the main driver of duration dependence in unemployment hazards and of the persistent scarring effects from nonemployment. Last, Grigsby (2022) focuses on skill heterogeneity as a key source of worker heterogeneity, developing a tractable, frictionless model of labor supply with multidimensional skills, and shows that composition and reallocation effects of non-uniform labor demand shocks can generate negative comovements between aggregate

[4] In Hornstein et al. (2011), all job-to-job transitions are wage- or value-improving, implying a high on-the-job offer arrival rate that lowers the reservation wage and rapidly moves workers up the job ladder. To account for wage cuts or job-to-job transitions that do not improve worker value, the model requires substantial dispersion in offered wages, thereby increasing the role of job heterogeneity.

wages and employment. Considering that effects of skills on workers' labor market outcomes are jointly determined by both the demand and supply of labor, the importance of skill aligns with the finding of Ahn (2023), which highlights the time-varying distribution of worker heterogeneity and its effects in unemployment hazards. Relatedly, Doniger (2023b) sheds light on the importance of contract heterogeneity as a source of wage differentials.

Last, this paper also relates to the empirical literature on on-the-job search driven by job-loss risks in the context of godfather shocks. Specifically, Fujita (2010) find, based on the UK Labour Force Survey, that workers search on the job for very different reasons, and that a nontrivial fraction engage in on-the-job search due to fear of losing their job. Similarly, Ahn and Shao (2021) provide empirical evidence, based on the American Time Use Survey, that workers are more likely to engage in on-the-job search, and spend more time on it conditional on searching, when the likelihood of job loss or nonemployment increases, suggesting a precautionary motive for on-the-job search. Simmons (2024) show similar findings based on a wider range of data including EU Labour Force Survey, British Household Panel Survey, UK Labour Force Survey, and Displaced Worker Supplement of the CPS.

3 Data

The Bureau of Labor Statistics' Survey of Income and Program Participation is comprised of a series of moderately length panels. In this study we employ the 1996, 2001, 2004, and 2008 panels, which share a common questionnaire and cover the time period 1996-2013. Surveys are conducted every four months and record, retrospectively, the experience of the respondent during the entire four month period. These records include weekly labor force status; labor earnings and weekly hours for each month; start and end dates of employment in up to two jobs per four month period; and the reason for separation for any job that ended in the four month period.

From these data we construct two data frames. First, a month-level record of labor force status and job transition for each respondent. This is used to estimate the MLMM described in Section 4. In this data frame, we record labor force status (employed, unemployed, inactive, missing); and job transition speed (same job, new job from 5 or more weeks of non-employment, new job from less than 5 weeks of non-employment, new job with less than 5 weeks overlap

Table 1: EMISSIONS OF MLMM

Panel A: Mean Observation

Labor Force Status		Job Transition Duration	
Employed	65.20	Same Job	58.63
Unemployed	4.98	New Job: > 5 weeks non-emp.	0.82
Non-participant	26.83	New Job: 0-5 weeks non-emp.	0.23
Non-response	2.99	New Job: 0-5 weeks overlap	0.12
		New Job: > 5 weeks overlap	0.03
		End of previous job, Non-response	40.17

Panel B: Other distributions

Number of Jobs		Job with Highest Hours		Job Transition Reason	
Zero	36.80	Full-time	45.22	Same Job	58.63
One	54.65	Part-time	13.93	> 5 weeks unemp.	0.82
Two	5.56	None	37.85	Involuntary separation	0.11
Non-response	2.99	Non-response	2.99	Job-to-job quit	0.21
				Other quit	0.06
				Non-response	40.17

Note: Authors' calculation.

with an old job, new job with more than 5 weeks of overlap with an old job, missing). The job transitions record the duration at transition and reason for transition to the job that is new in that month if there are two jobs within the two months prior to and including the month of transition. The respondent-level time-series of these states can be considered as emissions from the underlying hidden labor force state.[5]

Table 1 records the share of the data exhibiting each type of emission. Panel A records the emissions targeted in our MLMM estimation. Labor force status shares match the shares recovered from other surveys. While we do not force the transition data to be consistent with the labor force status data at the individual level, we see that, on average the vast majority of employed workers have tenure of at least one month (same Job). Meanwhile, the newly hired are predominantly individuals who were previously unemployed at least five weeks, with a short period of non-employment being the second most likely. Panel B records additional emissions

[5]It is important to note that the two sets of emissions satisfy the condition of conditional independence, an assumption required for the MLMM, which will be discussed in detail in the next section.

Table 2: SUMMARY STATISTICS: REGRESSION SAMPLE

ln(wages)	2.75
Hourly	*2.53*
Salary	*3.06*
Percent Newly Hired	22.96
NH&EE (percent of NH)	*18.03*
NH&EE&Separated (Percent of NH&EE)	*25.36*
NH&EE&JtJ Quit (Percent of NH&EE)	*61.11*
Percent Completed Tenure > 1 year	
Yes	87.99
No	12.00
Tenure censored	0.02
Age	39.61
Share Male	50.20
Education (years)	13.39
Potential Experience	20.22
Tenure	7.06

Note: Authors' calculation

that we can use to understand the MLMM result and/or as additional channels in alternative models. These show that the majority of job holders have exactly one job at a time and also that the majority of job holders hold full-time jobs. In addition, roughly fifty five percent of rapid job movers move due to a job-to-job quit while involuntary separations account for another thirty percent and other quits the remainder.

Our second data frame consists of a survey-wave×job-level record of real wages, labor force status, job transitions and demographics. This is used to estimate the wage regressions described in Section 6. This data frame differs from the MLMM data frame in important ways. Firstly, we include wage information. Real wages are calculated from job-specific monthly income, hours, and dates of employment and are deflated by the PPI. Second, we include up to two observations per respondent×time period, one for each possible job. Each is assigned the wage associated with that job-month. Third, we include only the month immediately preceding the SIPP interview. This selection renders us robust to the well documented "seam" bias.[6]

[6]Respondents have been documented to smooth wage information within 4-month survey waves more than

Table 2 records summary statistics for this second, regression analysis, data frame. As expected hourly workers earn a bit less than salaried. Approximately twenty three percent of workers are new hires in any given year. Of these, just less than 20 percent are hired within five weeks of the end of their previous employment spell. Of these, just over sixty percent report a job-to-job quit leading up to the latest accession and another twenty five percent report an involuntary separation (the remainder of the data is comprised of other quits and accessions with missing information about the end of the preceding job). Meanwhile, nearly ninety percent of job-month observations represent jobs that ultimately last at least one year. Age, sex, education, experience, and tenure data are consistent with other surveys.

4 Empirical Method

This section discusses the model in Section 4.1, the maximum likelihood estimation in Section 4.2, and the estimation result in Section 4.4.

4.1 Multi-channel Latent Markov Model

The goal is to uncover an individual's latent labor-market state at each point in time from the joint histories of reported labor force status and job-transition status observed in the SIPP. For this, we employ a multi-channel latent Markov model.[7]

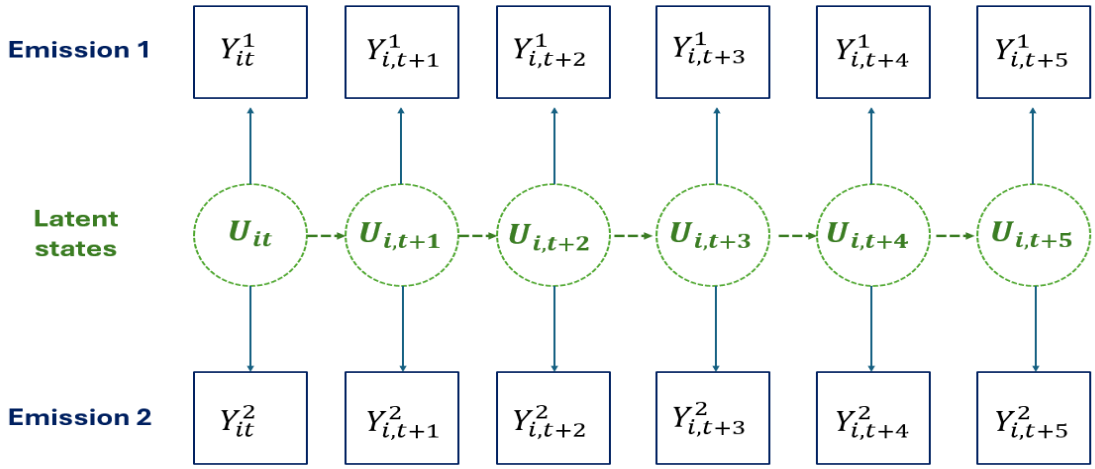
The MLMM consists of two components: the transition model and the emission model. The observed statuses are referred to as *emissions*, as they represent the signals individuals send regarding the latent state they occupy (Ahn et al., 2023).[8] The transition model describes the dynamics of latent states as a Markovian process. With the initial shares of latent states, the

across 4-month survey waves. Following the literature using SIPP wage information we assume that smoothing results from biased recall and note that such smoothing would attenuate coefficients in a regression of wages on time-varying indicators, such as the business cycle if each month were considered.

[7]We use the notation and the basic set-up of generalized latent Markov model in Bartolucci, Farcomeni and Pennoni (2013) and adapt to our case.

[8]This structure is similar to that of a state space model, which is composed of a state equation and a measurement equation. The state equation is analogous to the transition model, while the measurement equation corresponds to the emission model.

Figure 1: GRAPHICAL REPRESENTATION OF MLMM



Notes to figure: The figures illustrate the structure of the MLMM. The square boxes represent the emissions, with the upper and lower lines indicating the first and second emissions, respectively, from which the latent states – depicted as circles – are inferred. The horizontal arrows represent the Markovian transitions, while the vertical arrows indicate the emission process.
Source: Authors' illustration, adapted from Ahn et al. (2023).

transition model predicts the fraction of individuals in each latent state and the worker flows between latent states. The emission model connects these latent states to the observed labor force status and job status.

Figure 1 illustrates the basic structure of the MLMM. The circles represent latent states, with the horizontal arrows capturing the Markovian latent dynamics. The squares represent observed statuses. The vertical arrows linking the circles and squares depict the emission model. This model is described by three sets of parameters: (1) the initial probabilities of the latent states, (2) the Markovian transition probability matrix, and (3) the emission probabilities.

To explain the methodology intuitively, we consider a model for an individual i. Later on, all the individuals' cases are considered and integrated in the log likelihood function. For the ease of communication, we suppress the notation i. In addition, we use capital letters to denote random variables, and small letters to denote realizations of random variables. The bold notation is a vector, while the plain notation is a scalar.

Let $\mathbf{Y}_t$ denote the vector of two categorical response variables ($r = 2$) at time t. The first variable, denoted as Y^1_t, represents the reported labor force status. The second variable, denoted

as Y_t^2, represents the reported job-switch status. There are four possible categories for Y_t^1 ($c_1 = 4$), and seven possible categories for Y_t^2 ($c_2 = 7$). We stack $\mathbf{Y}_t$ for $t = 1, \ldots, T$ into the vector $\mathbf{Y}$ composed of $r\,T$ elements.

Let U_t represent the latent state at time t, where there are k possible states. By stacking the history of U_t, we obtain $\mathbf{U} = [U_1, U_2, \ldots, U_T]'$. The latent process is assumed to follow a first-order Markov chain.

The response vectors $\mathbf{Y}_1, \mathbf{Y}_2, \ldots, \mathbf{Y}_T$ are assumed to be conditionally independent or locally independent given the latent process $\mathbf{U}$. This assumption, known as *conditional independence* or *local independence*, is standard in the literature of latent Markov models. Specifically, Y_t^1 and Y_t^2, for $t = 1, 2, \ldots, T$, are conditionally independent from each other given U_t.

The emission model for the labor force status (the first emission model, henceforth) describes the probability that the individual reports a labor force status (y_1) out of E, U, N, and M conditional on the individual being in a particular latent state out of k-possible states (u). This probability, denoted ϕ_t^1 is written as follows:

$$\phi_u^1 = P(Y_t^1 = y_1 | U_t = u), \quad y_1 \in [E, U, N, M], \quad u \in [u_1, u_2, \cdots, u_k],$$

where u_p denotes p'th latent state for $p = 1, 2, \cdots k$.

Likewise, the emission model for the job-switch status (the second emission model, henceforth) describes the probability that the individual reports a job status (y_2) out of possible seven categories conditional on the individual being in a particular latent state out of k-possible states (u). This second emission probability, denoted ϕ_t^2, is written as follows.

$$\phi_u^2 = P(Y_t^2 = y_2 | U_t = u), \quad y_2 \in [j_1, j_2, \cdots, j_7], \quad u \in [u_1, u_2, \cdots, u_k]$$

where j_q denotes q'th job-switch status for $q = 1, 2, \cdots 7$.

The probabilities that the individual is in particular latent states—the initial probabilities—

are denoted as π_1 and written as the following.

$$\pi_1 = P(U_1 = u) \quad u \in [u_1, u_2, \cdots, u_k]$$

The probabilities that the individual switches from one latent state to another—the transition probabilities—are denoted as Π and described as

$$\Pi = P(U_t = u | U_{t-1} = u') \quad t = 2, 3, ..., T, \quad u', u \in [u_1, u_2, \cdots, u_k].$$

where u' is a realization of U_{t-1} that may be different from u, the realization of U_t.

Given the parameters, the probability distribution of entire history of latent states is expressed into the following.

$$P(\mathbf{U} = \mathbf{u}) = \pi_1 \prod_{t=2}^{T} \Pi_t$$

where $\mathbf{u}$ is the history of realized latent states spanning for $t = 1, 2, \cdots, T$.

The conditional distribution of $\mathbf{Y}$ given $\mathbf{U}$ may be expressed as

$$P(\mathbf{Y} = \mathbf{y} | \mathbf{U} = \mathbf{u}) = \prod_{t=1}^{T} \phi_t.$$

where we define

$$\phi_t = \prod_{j=1}^{2} \phi_t^j,$$

thanks to the conditional independence assumption on both emissions.

It is important to note that the two sets of emissions satisfy the condition of conditional independence. This condition means that, conditional on an individual's latent state, the reported labor force status and job status are independent of past emissions. To illustrate this in the context of labor force status: conditional on an individual's attributes and latent state, the person's report of labor force status in one month does not predict their response in the next month. The same logic applies to job status. Conditional on individual attributes, whether a

person is recorded to stay in the same job does not predict whether they will be recorded to have a new job or to be missing in the next period. Likewise, if a person is recorded to have a new job following a short unemployment spell, this does not predict whether they will later transition to a new job through unemployment or job overlap. We treat the two emissions as independent of each other and do not impose internal consistency between them. If there are inconsistencies – for example, if a respondent was employed last month but is now observed to have a new job following an unemployment spell longer than five weeks – the model uses such inconsistencies as informative signals when identifying latent states.

The realized distribution of observed responses about the labor force status and job-switch status is expressed as the following:

$$P(\mathbf{y}) = P(\mathbf{Y} = \mathbf{y}) = \sum_{\mathbf{u}} \pi_1 \times \Pi_2 ... \Pi_T \times \phi_1 \phi_2 ... \phi_T. \tag{4.1}$$

The computation of equation (4.1) involves summing over all possible configurations of latent states, totaling k^T combinations. We provide further discussions about the algorithm that is designed to mitigate the computational complexity in the next section.

So far, we have considered a case of the individual i for the ease of communication. In the estimation of the model, all the individuals in the sample are included in the likelihood function.

4.2 Maximum Likelihood Estimation

For an observed sample of N individuals, let $\mathbf{y}_i$ denote the history of i's responses. Since each individual provides two statuses at time t, $\mathbf{y}_i$ is composed of the sub-vectors $\mathbf{y}_{it} = [y_{it}^1, y_{it}^2]'$.

Assuming the independence among the individuals, the log-likelihood of the model is written as:

$$\ell(\theta) = \sum_{i=1}^{N} log P(\mathbf{y}_i), \tag{4.2}$$

where θ is the vector of all model parameters affecting $P(\mathbf{y}_i)$.

Equation (4.2) is rearranged into the equivalent form:

$$\ell(\theta) = \sum_{\mathbf{y}} n_{\mathbf{y}} log P(\mathbf{y}) \tag{4.3}$$

where $n_{\mathbf{y}}$ is the number of observation of a particular realized response sequence $\mathbf{y}$ in the sample. Computing the log likelihood function using equation (4.3) is more efficient than using equation (4.2), because the sum is restricted to the response sequences realized at least once and hence is observed in the sample.

The estimation process involves the Baum-Welch algorithm, also known as the forward-backward algorithm, for inferring latent states and the EM algorithm for parameter estimation. The estimation of the model involves considering all possible histories of latent states, which is numerically challenging. The Baum-Welch algorithm alleviates the computational difficulty: Leveraging the Markovian dynamics of latent states, a history of latent states is decomposed into the small portions that are easier to keep track of.[9]

The likelihood function is maximized using the EM algorithm, which helps to overcome the numerical difficulties associated with estimating a latent state model with a large number of parameters. The EM algorithm alternates two steps in the maximum-likelihood estimation. In the E-step, the sequences of latent states are predicted and updated from the previous iteration given the model parameters. In the M-step, the model parameters are re-estimated given the predicted sequences of latent states from the E-step. By separating these two steps, the optimization is faster than the case of one-step estimation. Note that the Baum-Welch algorithm is employed in the E-step.

Last, using the Viterbi algorithm (Viterbi, 1967; Juang and Rabiner, 1991) and the estimated model parameters, we also recover the most likely individual sequence of latent states.[10] This technique allows us directly observe the full history of experiencing latent states over the sample

[9]The Baum-Welch algorithm also produces elements used to produce the locally and globally decoded sequences of latent states at the individual level.

[10]In the literature of hidden Markov model, this procedure is also known as *"decoding."* Additional discussions on decoding are in the appendix.

period for each individual. Further details are found in the appendix.

4.3 Number of latent states and Identification of MLMM

4.3.1 Number of latent states

So far, we have not determined the number of latent states (k). Theorem 6 of Allman, Matias and Rhodes (2009) provides a condition about identifiable number of latent states, which we discuss below.

The parameters of a latent Markov model with k hidden states and N observable states are identifiable from the marginal distribution of $T = 2q + 1$ consecutive variables provided q satisfies the following condition:

$$\binom{q+N+1}{N-1} \leq k.$$

This result indicates that with four labor force statuses ($N = 4$) and a total of 64 months of observations ($T = 64$), up to 7,140 latent states can be identified. With the second emission, which includes seven categorical responses, it is theoretically possible to identify more than 7,140 latent states.

Meanwhile, determining how many latent states the data can support in an empirically distinct and meaningful way is a matter of model selection. In this context, we use both the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) to guide our choice of the number of states.[11] Given the likelihood estimates, we can choose the optimal number of latent states based on formal statistical criteria, in this way minimizing ad-hocness in our statistical analyses.[12]

[11] The BIC is usually preferred if the two criteria do not agree. See for example Bartolucci et al. (2013) and Bacci, Pandolfi and Pennoni (2014a).

[12] The formulation of AIC and BIC is in the appendix.

4.3.2 Discussion on Model Identification

Identification of the parameters in a multi-channel latent Markov model with two emissions requires that the joint distribution of the observables uniquely determine the model parameters, up to a relabeling of the latent states.[13] The identification of MLMM essentially builds upon the general identification condition of a Markov model, and is discussed in the identifiability results for finite mixtures and hidden Markov models from Allman et al. (2009).

To demonstrate that the model is identified, we outline the identification conditions and relate them to the estimation results presented in the next section.

First, the conditional independence, an identifying assumption, should be satisfied. This restriction allows information from each channel to contribute separately to state discrimination, improving identification. As previously discussed in Section 3, we construct our dataset in a way that satisfies the conditional independence.

Second, the emission distributions $\{\phi^1_{u_p}(\cdot), \phi^2_{u_p}(\cdot)\}^k_{p=1}$ for a given latent state u_p must be sufficiently distinct across latent states. If two states have identical emission distributions for both channels, they are observationally indistinguishable and cannot be separately identified. As shown in Table 4.4, the emission probabilities are distinct across the seven latent states.

Third, related to the first condition and Allman et al.'s theorem, the emission distributions should be non-degenerate and contain richness of information enough to distinguish k number of latent states. In this sense, multiple channels can enhance identifiability by increasing the number of distinct observable distributions linked to each state. We demonstrate this condition has been satisfied, as discussed above.

Fourth, related to the first three points, identification of latent states also relies on variations in observed sequences across time. In this sense, longer panels provide more transitions and improve distinction of the transition dynamics from the emission parameters. Our dataset covers total 64 monthly observations for a given individual, which is sufficiently long enough to

[13]The likelihood is invariant to permutations of the state labels. Identification is therefore only up to label switching, which must be resolved by normalizations or post-estimation constraints. We get around with this issue by imposing restrictions on the relative magnitudes on the transition and emission probabilities of each latent state.

Table 3: ESTIMATES OF INITIAL SHARES AND EMISSION PROBABILITIES.

	1 (M)	2 (Em)	3 (Es)	4 (N)	5 (U)	6 (Xn)	7 (Xs)
Initial Share	0.08	0.07	0.52	0.25	0.05	0.01	0.02
Emission 1	1	2	3	4	5	6	7
NA	1.00	0.00	0.00	0.00	0.00	0.00	0.00
Employment	0.00	1.00	1.00	0.00	0.00	0.46	0.23
Unemployment	0.00	0.00	0.00	0.00	0.98	0.26	0.47
NLF	0.00	0.00	0.00	1.00	0.01	0.28	0.30
Emission 2	1	2	3	4	5	6	7
NA	1.00	1.00	0.00	1.00	1.00	0.04	0.01
Same job	0.00	0.00	0.99	0.00	0.00	0.03	0.94
New job and							
non-emp > 5 weeks	0.00	0.00	0.00	0.00	0.00	0.67	0.00
non-emp 0-5 weeks	0.00	0.00	0.00	0.00	0.00	0.00	0.04
0-5 weeks overlap	0.00	0.00	0.00	0.00	0.00	0.00	0.00
> 5 weeks overlap	0.00	0.00	0.00	0.00	0.00	0.01	0.00
end of previous job, Non-response	0.00	0.00	0.00	0.00	0.00	0.24	0.00

Source: Survey of Income and Program Participation and author's calculations.

Note: This table reports the estimates of initial probabilities and emission probabilities. NA denotes not available or missing, and NLF denotes not in the labor force.

precisely identify the latent states.

Last, the transition probability matrix must be full rank and ergodic. This condition ensures that each latent state has Markovian dynamics that distinguishes one from another. The determinant of the estimated transition probability matrix is non-zero.[14]

4.4 Estimation Results

Table 3 presents the estimated initial share of each state and the emission probabilities for each state. Table 4 reports the estimated Markovian transition probabilities. The seven latent states, unevenly distributed across the population, exhibit distinct emissions and latent dynamics.

As summarized in Table 3, the seven latent states are distinguished by their emission probabilities. Based on their first emission of labor force status (Emission 1), states 2 and 3 are E-prone, state 4 is N-prone, state 1 is M-prone, and state 5 is U-prone. States 6 and 7 exhibit non-negligible emission probabilities across all three labor force statuses, and we refer to them as "X-prone." According to the initial share, the E-prone and N-prone states account for approximately 60 percent and 25 percent of the population, respectively. The M-prone state comprises 8 percent, the U-prone state 5 percent, and the two X-prone states together make up only 3 percent.

The two E-prone states and the two X-prone states are further distinguished by their emissions about job information (Emission 2). Most individuals in the E-prone states report having the *same* job as in the previous month (state 3), representing stable employment. Meanwhile, 10 percent of those in the E-prone category have their job information *missing* (state 2) – likely reflecting those who reported being employed but refused questions about their employer. We call the first E-prone state as *Es*-prone and the second as *Em*-prone. The two X-prone states are distinguished by their likelihoods of staying employed as well as whether they have a new job or stay at the same job. Individuals in the first X-prone state (state 6) are employed about half the time, higher than those in the second X-prone state (state 7) who are employed roughly a quarter of the time. The first X-prone state —comprising only one percent of the sample— is likely associated with accessing to a *new* job, and the second X-prone state, only two percent

[14] We explicitly account for non-responses to address potential non-randomness in missing observations. The missing state exhibits high persistence, indicating that most missing observations correspond to individuals who have permanently left the sample. In this sense, the missing state functions like an absorbing state, which makes the ergodicity condition somewhat fragile. As an alternative, we estimate a model that does not treat non-responses as a separate category, yielding transition probabilities consistent with an ergodic Markov chain. The estimation results remain robust under this specification, although state 1 no longer appears.

Table 4: ESTIMATES OF MARKOVIAN TRANSITION PROBABILITIES.

		To state …						
		1 (M)	2 (Em)	3 (Es)	4 (N)	5 (U)	6 (Xn)	7 (Xs)
From state …	1	**0.99**	0.00	0.00	0.00	0.00	0.00	0.00
	2	0.04	**0.93**	0.01	0.01	0.00	0.01	0.00
	3	0.04	0.00	**0.94**	0.00	0.00	0.00	0.02
	4	0.04	0.00	0.01	**0.92**	0.01	0.02	0.00
	5	0.04	0.01	0.01	0.08	**0.75**	0.10	0.01
	6	0.04	0.01	0.78	0.01	0.01	**0.01**	0.15
	7	0.05	0.01	0.18	0.16	0.12	0.03	**0.45**

Source: Survey of Income and Program Participation and author's calculations.

Note: This table reports the probabilities of switching between the seven latent states.

of the sample, is associated with having the *same* job. We call the first X-prone state as *Xn*-prone and the second X-prone state as *Xs*-prone. These patterns suggest that the circumstances surrounding job switches provide identifying information about latent labor-market states.

The seven latent states exhibit distinct levels of persistence and transition dynamics, as summarized in Table 4. The diagonal elements of the transition probability matrix reflect the persistence of each state, while the off-diagonal elements indicate the likelihood of transitioning to other states. The states prone to U, *Xn*, and *Xs* are noticeably transient, with continuation probabilities of 0.75 and below 0.5, respectively – compared to over 0.9 for the more stable states.

In particular, the *Xn*-prone state has an especially low continuation probability of just one percent, making it the *most transient* state. Nearly all individuals in this state move to a different state in the next month: About three-quarters transition into stable employment, while a non-negligible share – around 10 percent – move into the *Xs*-prone state.

Importantly, the three transient states are interconnected. Approximately 10 percent of individuals in the U-prone state (state 5) transition to the *Xn*-prone state (state 6) – the mixed state involving new job starts after long jobless spells – and then to the *Xs*-prone state (state 7). The *Xs*-prone state is marked by frequent non-employment but some employment stability. From there, some individuals return to the U-prone or non-employment-prone (N-prone) states.

Additionally, a small share of individuals initially in the *Es*-prone state (state 3) move to state the *Xs*-prone state, then to the *U*-prone state (state 5), and subsequently to the *Xn*-prone state (state 6) – the most transient state. Essentially, these flows capture paths for workers who lose jobs, experience non-employment, secure new employment temporarily, but fail to transition into stable work.

Figure 2 shows the latent state paths of selected individuals, illustrating stability and vulnerability, respectively. In the top panel, two individuals remain in the *Es*-prone state (state 3) associated with stable employment – in a highly persistent manner. In contrast, the bottom panels display two individuals who make frequent transitions across states. Notably, despite being randomly selected, all three of these individuals frequently pass through the transient states (states 5,6, and 7).

These dynamics highlight that the two mixed states, together with the U-prone state, represent volatile labor market experiences. Affecting only a small share of the population, these states are likely associated with the pockets of instability and vulnerability in the labor market.[15] We discuss related details, including latent types, in the next section.

5 Identification of Vulnerable Type

Section 5.1 introduces the identification of worker types, and Section 5.2 discusses the key features of the estimated types. Section 5.3 examines the observed attributes of individual types and explores their implications for structural modeling.

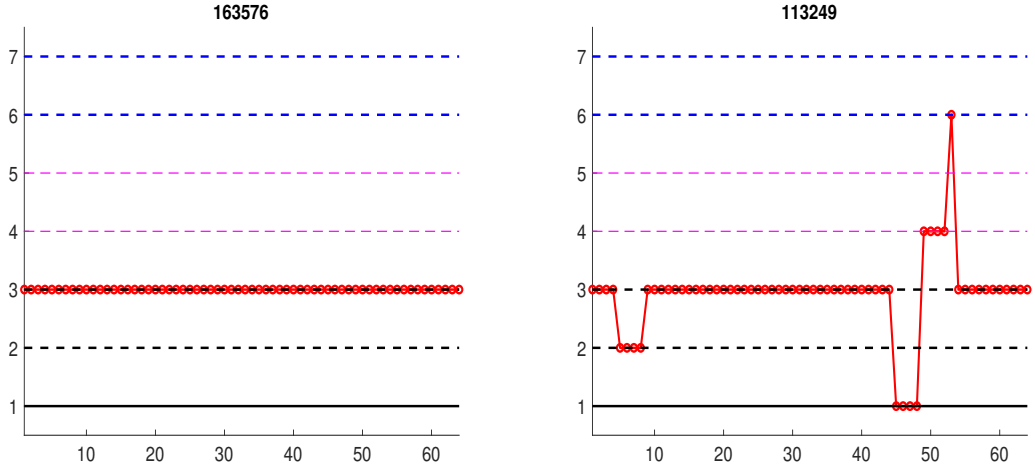
5.1 Method

We classify a worker's type based on their history of latent states, particularly their experience of the transient states – the *U*, *Xn*, and *Xn* -prone states. Specifically, if an individual is in one of the transient states in month t, we classify them as a vulnerable type from *"b"* months before to

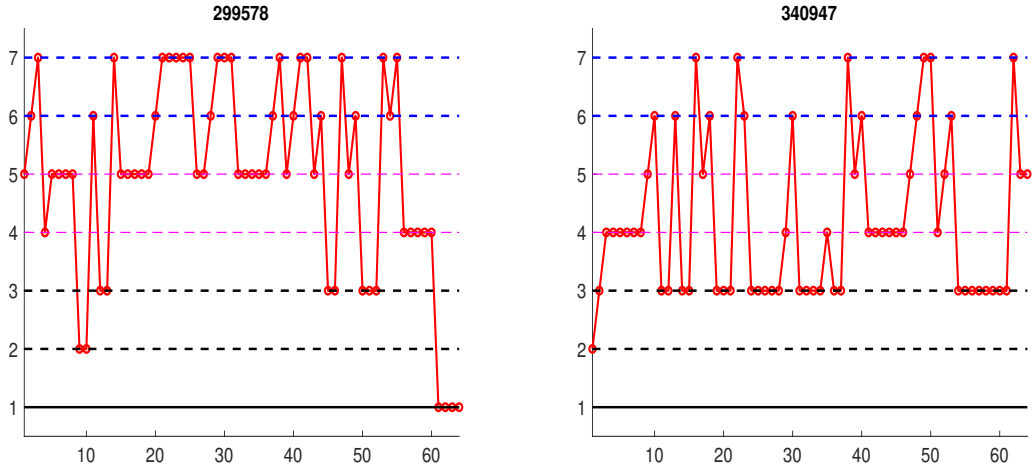
[15] In the appendix, we present the relationship between latent states and observable worker attributes, such as demographics, industry, and occupation.

Figure 2: LATENT STATE HISTORIES OF EMPLOYMENT-PRONE WORKERS

(Stable Histories)



(Transient Histories)



Note: This figure shows the latent state histories for selected individuals from the sample. The number at the top of each panel represents the individual identifier in our dataset (Identifier to be updated). The Y-axis indicates state numbers, while the X-axis represents the months in the sample. The black dashed lines capture the employment-prone states, the magenta dashed lines capture non-employment prone states, and the blue dashed lines capture transient states.
Source: Authors' calculations.

"b" months after t, covering a total of $(2b+1)$ months. Because individuals often move between the transient states, as observed in Figure 2, the actual duration of being classified as vulnerable typically may exceed $2b+1$ months. We call "b" as the bandwidth.[16] The bandwidth size captures information related to a specific degree of persistence. A larger bandwidth corresponds to more

[16]The intuition behind this approach is similar to that of the bandwidth filter.

persistent types, but also increases the likelihood of a Type II error (false positive). Conversely, a smaller bandwidth identifies less persistent types, which may be conceptually less consistent with the idea of temporal stability, but reduces the risk of a Type II error.

We select the bandwidth based on two considerations.

- **[External validation]** The identified type provides reasonable predictive power for variables not used in its classification when included in the prediction model..
- **[Internal consistency]** The ex-post distribution of type durations and the probabilities of type continuation are consistent with the chosen bandwidth.

For the first criterion, we examine the coefficient on unemployment in the wage regression that includes individual type. Since wage changes are generally acyclical or procyclical, we do not expect a statistically significant positive coefficient. A positive relationship would suggest that workers become more productive or demand higher wages as unemployment rises– implying countercyclical wage changes even after controlling for individual type. This would indicate that the type absorbs substantial procyclical variation, allowing real wages to vary countercyclical conditional on the type. Such a pattern conflicts with the notion of type as a temporally stable characteristic. Therefore, we consider the type to be appropriately identified when the unemployment coefficient is either statistically insignificant or negative.

It is important to highlight two compelling features of our type classification based on latent states. First, the classification leverages rich information from individuals' histories of labor force status, job transitions, and the context of those transitions, as captured by the estimated latent states and their Markovian dynamics. Second, the classification allows individual types to evolve gradually over time. This approach ensures temporal stability –consistent with the concept of worker types – while allowing sufficient flexibility for type changes over the relatively long duration of individual histories in the SIPP. This dynamic feature distinguishes our approach from previous studies (e.g., Gregory et al., 2025; Borovičková and Macaluso, 2025) that rely on K-means clustering and assume fixed types over time.

Table 5: SUMMARY OF STABLE AND VULNERABLE TYPES

	Stable	Vulnerable
Share (%)	80	20
Type continuation probability (without missing)	0.99	0.94
Average duration (months)	44.5	15.5

Source: Authors' calculation.

5.2 Estimation Results

We select a bandwidth of $b = 5$ with a type duration of 11 months. This bandwidth satisfies the aforementioned model selection criteria. It is important to note that the duration of each type is long enough to distinguish workers with seasonal employment patterns from those without seasonal employment risk. A longer bandwidth increases the risk of type-II error and fails to satisfy external validity, while a shorter bandwidth cannot accommodate seasonal turnover, leading to artificial seasonality in type changes that is inconsistent with the concept of stable individual types.

Table 5 reports the share and characteristics of each identified type over the period 1996–2021. Overall, the stable type accounts for approximately 80 percent of the sample, while the vulnerable type comprises about 20 percent.[17] Second, the probability of remaining in the same type conditional on having been in that type in the previous month is 0.99 for the stable type and 0.95 for the vulnerable type. These high continuation probabilities are consistent with the interpretation of type as capturing temporal stability. Third, the average monthly duration is 44.5 months for the stable type and 15.5 months for the vulnerable type.

5.3 Observed Attributes of Types

Since we directly observe each individual's type over time, we can examine the composition of demographic attributes, job characteristics, industries, and occupations within each type.

Table 6 summarizes the socio-economic attributes of workers by type. Type 0 and Type 1 refer

[17] This distribution broadly aligns with the population share of the secondary labor market identified by Ahn et al. (2023).

Table 6: DEMOGRAPHIC COMPOSITION OF THE TYPES

	Type=0 (Stable)	Type=1 (Vulnerable)
Age	40.5	35.4
Education		
Less than High school	0.07	0.13
High school / some college	0.62	0.68
College or more	0.31	0.20
Female	0.5	0.5
Tenure	9.2	3.3
Tenure on previous job if any	3.8	2.2
ln(real wage)	2.8	2.5
Potential experience	20.9	16.5
Number of jobs	1.1	1.1
Full time (if employed)	0.75	0.52
Paid hourly (if employed)	0.54	0.77
Unionized (if employed)	0.14	0.08
Separation rate (monthly)	0.020	0.109
Accession rate (monthly)	0.004	0.101

Note to figure: The sample is wage earners waged 20-60. **Source**: Authors' calculation.

to the stable and vulnerable types, respectively. The sample is restricted to wage earners aged 20-60. For ease of communication, we refer to Type 0 individuals as the stable type and Type 1 individuals as the vulnerable type. Compared to the vulnerable type, the stable type is older on average, with a mean age of 41 versus 35. The stable type also has slightly more education – an average difference of 0.7 years – though this gap is small. Both types include more men, reflecting their greater share in the labor force. Overall, demographic differences between the two types are limited.

The two types differ markedly in employment persistence and job turnover. The average tenure exceeds 9 years for the stable type, compared to just 3 years for the vulnerable type. Separation and accession rates for the vulnerable type are approximately five and twenty times higher, respectively, reflecting their lower job stability. The stable type also has longer potential experience, likely due to accumulated job-specific human capital that strengthens attachment to their current job. As a result, the stable type earns real wages about 30 percent higher than the vulnerable type.

Table 7: OCCUPATION COMPOSITION OF THE TYPES

	Type=0 (Stable)	Type=1 (Vulnerable)
Managerial and Professional	31.5	19.6
Technical, Sales, Administration	29.8	28.3
Service	14.7	21.9
Precision Production	9.6	10.5
Operators, Fabricator	12.9	18.6
Military	0.2	0.1
missing (%)	1.4	1.1

Note to figure: The sample is wage earners waged 20-60. **Source**: Authors' calculation.

Table 8: INDUSTRY COMPOSITION OF THE TYPES

	Type=0 (Stable)	Type=1 (Vulnerable)
Agriculture, Forestry	1.31	2.72
Mining	5.22	9.26
Nondurable Manufacturing	5.83	4.95
Durable Manufacturing	9.57	7.81
Transportation, Communication	7.2	5.43
Wholesale Trade	3.67	2.74
Retail Trade	14.87	21.01
Finance, Insurance, and Real Estate	6.65	4.29
Business and Repair Services	5.76	8.88
Personal Services	2.46	3.71
Entertainment and Recreation	1.37	2.28
Professional and Related Services	28.79	24.1
Public Administration	6.79	2.64
missing (%)	0.53	0.19

Note to figure: The sample is wage earners waged 20-60. **Source**: Authors' calculation.

Next, Tables 7–8 present the distribution of occupations and industries across types. The stable type is more likely to work in managerial and professional occupations, while the vulnerable type is more concentrated in service occupations and in roles such as operators and fabricators. In addition, the vulnerable type takes a larger share in agriculture, mining, retail trade, and a few service sectors such as business and repair services and personal services than the stable type. These industries are typically affected by seasonal demand fluctuations (e.g., agriculture, forestry, mining, and retail services) and include many low-skill service sectors.

The observable attributes of each type suggest that type classification is likely determined by the interaction between job-side demand and worker-side supply factors. The concentration of the vulnerable type in industries with seasonal demand and turnover implies exposure to seasonal risks. Workers whose seasonal jobs end may transition into low-skill service sectors, which are also highly sensitive to aggregate demand, thereby exposing them to cyclical risks as well. Frequent turnover and job switching may hinder the accumulation of job-specific human capital, making it more difficult for these workers to secure stable employment and build job-specific skills.

6 Latent Heterogeneity in Wage Changes of Job Switchers

Using the estimated individual histories of latent states and types, we examine the extent to which these factors predict individual wage changes and contribute to the excess cyclicality of wage changes among new hires.

Consider the regression model as follows:

$$\begin{aligned}\Delta \ln(real\ wage_{i,t}) = {} & \beta_0 + \beta_1 \triangle UR_t + \beta_2 T_{1t} + \beta_3 T_{2t} + \beta_4 (T_{1t} \times \triangle UR_t) + \beta_5 (T_{2t} \times \triangle UR_t) \\ & + \beta_6 NH_t + \beta_7 (NH_t \times \triangle UR_t) + \beta_8 (NH_t \times T_{1t}) + \beta_9 (NH_t \times T_{2t}) \\ & + \beta_{10} (NH_t \times T_{1t} \times \triangle UR_t) + \beta_{11} (NH_t \times T_{2t} \times \triangle UR_t) \\ & + controls_{i,t}\Gamma + \varepsilon_{i,t},\end{aligned}$$

where $\triangle UR_t$ is the first-differenced unemployment rate in month t, T_{1t} is the indicator of the vulnerable type in month t, T_{2t} is the indicator of the stable type in month t, and NH_t is the indicator of being a new hire in month t. The term $controls_{i,t}$ denotes individual-level controls, including education, experience and its square, tenure, time fixed effects, and a constant, and the notation Γ is the vector of coefficients on $controls_{i,t}$.[18]

Table 9 reports the regression results. Several key findings emerge. First, while accession to a new job (New hire) significantly raises wages, being of the vulnerable type (Vulnerable×New hire) does not, more than offsetting the gains from starting a new job. Second, real wage changes for vulnerable workers—both incumbents and new hires—are significantly negatively correlated with changes in the unemployment rate gap, in stark contrast to the acyclicality observed in typical real wage changes.[19]

We further distinguish between workers who make direct EE transitions and those who make EE transitions with an intervening nonemployment spell (ENE). Table 10 reports the regression results. Vulnerable-type ENE transitioners experience wage cuts, whereas vulnerable-type EE transitioners realize smaller wage gains than incumbents. Notably, wage changes for vulnerable-type EE transitioners exhibit statistically significant cyclicality.

Next, we compare workers with long tenure (more than one year in the previous job) and short tenure (less than one year). Table 11 reports the regression results. Vulnerable-type new hires with long tenure experience wage cuts, and the excess cyclicality of wage changes is primarily driven by this group. Longer jobless spells further exacerbate both the magnitude of wage cuts and the degree of excess cyclicality, as shown in Table 12.

This result has important implications for wage dynamics. Prior literature has shown that job-to-job transitions generally lead to wage gains relative to previous jobs, and that wages for

[18] Controls are a quadratic in tenure and calendar time (following Gertler et al. (2020)). For ease of exposition, coefficients are multiplied by 100. Standard errors are clustered at the individual and month. Data are the SIPP 1996-2008 panels, covering years 1996-2014. We include both sexes and workers paid hourly as well as salaried. $N = 1,072,751$ and $Adj.R^2 = 0.71$. GHT restrict to males paid hourly, our results are robust to these sample selections.

[19] We further investigate which latent states drive the negative impact of the vulnerable type on new hires' wages and find that the *Xn*- and *Xs*-prone states are the primary contributors to both the negative effect and the excess cyclicality.

Table 9: RESULT OF WAGE REGRESSION (1)

	(1)	(2)	(3)
Dependent variable	current month +/- 5 months ...		
ln(real wage)	Before (-5)	After (+5)	Symmetric
Change in UR gap	0.00781***	0.000120	0.00179
	(0.00202)	(0.00182)	(0.00187)
Vulnerable	-0.00157	0.00279*	0.00141
	(0.00142)	(0.00150)	(0.00116)
Vulnerable × Change in UR gap	-0.0146**	-0.00197	-0.00481
	(0.00648)	(0.00645)	(0.00512)
New hires	0.0639***	0.0744***	0.0730***
	(0.00297)	(0.00263)	(0.00289)
New hires × Change in UR gap	0.00180	0.00912	0.0142
	(0.0117)	(0.0103)	(0.0115)
Vulnerable × New hires	-0.0704***	-0.0909***	-0.0801***
	(0.00368)	(0.00347)	(0.00351)
Vulnerable × New hires × Change in UR gap	-0.0110	-0.0227*	-0.0270**
	(0.0135)	(0.0123)	(0.0127)
Observations	1,079,744	1,079,744	1,079,744
R^2	0.001	0.002	0.001

Note to table: Standard errors in parentheses
The coefficients on unassigned type are not reported.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

new hires tend to be acyclical, as job changes are often driven by worker preferences. In addition, workers with long tenure typically hold more stable jobs, and their wages are less sensitive to business cycle fluctuations. In contrast, our empirical findings reveal a distinct group of vulnerable workers whose wage dynamics differ markedly from these established patterns. First, these workers switch jobs even when doing so does not result in wage gains – and in many cases, leads to wage losses. Second, their wages are procyclical with statistical significance, in contrast to the acyclicality observed for most workers. This suggests that their job transitions may be involuntary and driven by cyclical risks. Third, long tenure at the previous job does not protect workers from vulnerability; rather, it tends to have a more negative effect on subsequent wages, likely reflecting both the destruction of job-specific human capital and the greater reallocation frictions faced by vulnerable workers with long tenure.

Table 10: RESULT OF WAGE REGRESSION (2)

	Before (-5m)	After (5m)	Symmetric
Change in UR gap	0.00775***	0.0000282	0.00170
	(0.00201)	(0.00181)	(0.00186)
Vulnerable	-0.00140	0.00288*	0.00157
	(0.00141)	(0.00149)	(0.00116)
Vulnerable × Change in UR gap	-0.0147**	-0.00199	-0.00485
	(0.00645)	(0.00642)	(0.00509)
ENE	0.0830***	0.0697***	0.0918***
	(0.0171)	(0.0155)	(0.0175)
EE	0.0691***	0.0796***	0.0799***
	(0.00370)	(0.00315)	(0.00351)
ENE × Change in UR gap	0.00688	0.0146	0.0167
	(0.0165)	(0.0152)	(0.0173)
EE × Change in UR gap	0.00737	0.000144	0.0171
	(0.0190)	(0.0154)	(0.0173)
Vulnerable × ENE	-0.115***	-0.105***	-0.125***
	(0.0173)	(0.0158)	(0.0177)
Vulnerable × EE	-0.0404***	-0.0660***	-0.0538***
	(0.00492)	(0.00466)	(0.00464)
Vulnerable × ENE × Change in UR gap	-0.0135	-0.0263	-0.0271
	(0.0178)	(0.0166)	(0.0181)
Vulnerable × EE × Change in UR gap	-0.0625***	-0.0482**	-0.0690***
	(0.0233)	(0.0207)	(0.0213)
Observations	1,070,149	1,070,149	1,070,149
R^2	0.001	0.002	0.001

Note to table: Standard errors in parentheses
The coefficients on unassigned type are not reported.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 11: RESULT OF WAGE REGRESSION (3)

	All	prior tenure<1	prior tenure≥1
Change in UR gap	0.00201	0.000989	0.00192
	(0.00180)	(0.0106)	(0.00179)
Vulnerable	-0.00174	0.00557	0.00147
	(0.00170)	(0.00366)	(0.00223)
Vulnerable × Change in UR gap	-0.0147*	-0.00936	-0.0180*
	(0.00786)	(0.0181)	(0.00968)
New hire	0.0681***	0.0957***	0.0558***
	(0.00273)	(0.00553)	(0.00343)
New hire × Change in UR gap	0.00524	0.00130	0.00923
	(0.0103)	(0.0202)	(0.0139)
New hire × Vulnerable	-0.0786***	-0.0633***	-0.121***
	(0.00364)	(0.00667)	(0.00486)
New hire × Vulnerable × Change in UR gap	-0.00807	-0.0133	-0.00230
	(0.0131)	(0.0253)	(0.0172)
Education years	-0.000734***	-0.00144**	-0.000586***
	(0.000186)	(0.000724)	(0.000191)
refdate	-0.00000161***	-0.00000219***	-0.00000150***
	(0.000000191)	(0.000000747)	(0.000000197)
Experience	-0.00145***	-0.00179***	-0.00131***
	(0.000128)	(0.000436)	(0.000136)
Experience × Experience	0.0000269***	0.0000325***	0.0000239***
	(0.00000298)	(0.0000118)	(0.00000312)
Tenure	-0.000259***	0.00198	-0.000221***
	(0.0000520)	(0.00623)	(0.0000526)
Constant	0.0613***	0.0688***	0.0566***
	(0.00402)	(0.0157)	(0.00415)
Observations	1,079,983	103,154	976,828
R^2	0.001	0.007	0.001

Note to table: Standard errors in parentheses
The coefficients on unassigned type are not reported.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 12: RESULT OF WAGE REGRESSION (4)

	All	prior tenure<1	prior tenure≥1
Change in UR gap	0.00193	0.000525	0.00189
	(0.00179)	(0.0105)	(0.00179)
Vulnerable	-0.00158	0.00611*	0.00154
	(0.00170)	(0.00362)	(0.00222)
Vulnerable × Change in UR gap	-0.0147*	-0.00932	-0.0181*
	(0.00782)	(0.0179)	(0.00964)
Vulnerable × Change in UR gap	-0.000754	-0.0287	0.00541
	(0.0191)	(0.0525)	(0.0207)
New hire, long gap	0.0846***	0.0744***	0.105***
	(0.0159)	(0.0265)	(0.0221)
New hire, short gap	0.0726***	0.0957***	0.0625***
	(0.00334)	(0.00637)	(0.00433)
New hire, long gap × Change in UR gap	-0.00103	0.00391	-0.00144
	(0.0147)	(0.0250)	(0.0217)
New hire, short gap × Change in UR gap	0.0194	-0.00102	0.0360*
	(0.0163)	(0.0312)	(0.0209)
New hire, long gap × Vulnerable	-0.116***	-0.0541**	-0.195***
	(0.0161)	(0.0268)	(0.0224)
New hire, short gap × Vulnerable	-0.0478***	-0.0438***	-0.0760***
	(0.00488)	(0.00828)	(0.00699)
New hire, long gap × Vulnerable × Change in UR gap	0.00107	-0.0141	0.0122
	(0.0167)	(0.0292)	(0.0240)
New hire, short gap × Vulnerable × Change in UR gap	-0.0792***	-0.0473	-0.107***
	(0.0220)	(0.0398)	(0.0294)
Education years	-0.000740***	-0.00152**	-0.000594***
	(0.000185)	(0.000729)	(0.000191)
refdate	-0.00000152***	-0.00000179**	-0.00000146***
	(0.000000191)	(0.000000756)	(0.000000196)
Experience	-0.00137***	-0.00178***	-0.00123***
	(0.000128)	(0.000440)	(0.000136)
Experience × potexp	0.0000251***	0.0000319***	0.0000222***
	(0.00000298)	(0.0000119)	(0.00000312)
Tenure	-0.000218***	0.00698	-0.000196***
	(0.0000518)	(0.00632)	(0.0000525)
Constant	0.0590***	0.0600***	0.0550***
	(0.00401)	(0.0158)	(0.00415)
Observations	1,070,348	99,129	971,218
R^2	0.001	0.007	0.001

Note to table: Standard errors in parentheses
The coefficients on unassigned type are not reported.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

7 Implications for theory

We explore the theoretical implications of our empirical findings using a quantitative on-the-job search model with the godfather shock. The "godfather shock" is an on-the-job offer that can not be refused, considered by Tjaden and Wellschmied (2014). A baseline on-the-job search model without such shocks implies 1) that all job-to-job transitions are value improving for the worker and 2) limited frictional wage dispersion (Hornstein et al., 2011). Tjaden and Wellschmied (2014) show that introducing godfather shocks provides an explanation for the observed incidence and magnitude of wage cuts between adjacent jobs and facilitates a calibration that features greater wage dispersion. We use the shares of workers experiencing transient states, the vulnerable type, and changes in the incidence of transient states, we infer the share affected by godfather shocks and show that the quantitative on-the-job search model replicates our key empirical patterns—wage-degrading EE transitions and excess wage cyclicality among new hires of the vulnerable type.

To grasp the intuition, note that the observed job-to-job transition incidence in a model with godfather shocks will be the sum of the incidence of transitions that reveal preference and the incidence of offers that can not be refused:

$$JTJ = \underbrace{\lambda(1-\lambda_d)\int_{w^*}^{w_{max}}[1-F(z)]dG(z)}_{\text{incidence of preference revealing moves}} + \underbrace{\lambda\lambda_d[1-F(w^*)]}_{\substack{\text{incidence of offers that} \\ \text{can not be refused}}}, \tag{7.1}$$

where λ is the hazard of on-the-job offers, λ_d is the probability that an offer is a godfather shock, and $F(z)$ and $G(z)$ are the distributions of wage offers and of wages of employees in equilibrium. $F(z)$ and $G(z)$ have support $[w^*, w^{max}]$ where w^* is the lowest reservation wage across worker types and w^{max} is determined by the distribution of firms' marginal product of labor together with job finding and separation hazards. Observe that in the case of a preference revealing transition, the wage of the firm to which the worker transitions is a function of the current incumbent's wage. Specifically, the new wage must be in the set $(z, w^{max}]$. Meanwhile, in the

Table 13: SHARE OF JOB-TO-JOB (JTJ) TRANSITIONS FROM THE TRANSIENT STATE

Transition	Share of Job-to-Job transitions
State 5 → State 3	≈ 0.12
State 6 → State 3	≈ 0.32
State 7 → State 3	≈ 0.36
States 5/6/7 → State 3	≈ 0.51

Note to figure: This table shows the share of EE transitioners who passed through the transient latent states. States 5, 6, and 7 are classified as transient, while State 3 represents stable employment, as identified by the MLMM.
Source: Authors' calculation

case of the godfather shock the new wage is drawn from the full support, that is $[w^*, w^{max}]$. Tjaden and Wellschmied (2014) and subsequent papers identify λ_d jointly from the observed job-to-job transition rate and the incidence of job-to-job transitions occurring at a wage cut.

Here we take a different approach and link the rate of godfather shocks the share of observed job-to-job mobility that occurs whilst in the vulnerable state. While we acknowledge that the mapping may be imperfect, this strategy has the key advantage of identifying the model while leaving the wage distribution and wage-change distributions free for testing model implications. Before proceeding, we document that our measure aligns with that of Tjaden and Wellschmied (2014).

In the previous section, we documented higher separation rates for workers in the transient states. Table 13 summarizes the accession rates by previous transient state. Note that the accession from States 5, 6, and 7 to State 3 (last row of the table) represents the rate at which vulnerable-type workers move into stable employment. Using each of these values as the share of workers exposed to the godfather shock, we estimate the corresponding mean and median changes.

Table 14 reports the mean and median wage changes for job-to-job transitioners who were in transient states prior to their transitions. Overall, the model yields lower wage changes for vulnerable-type workers compared with stable types. Notably, job changers who experienced unemployment between two employment spells incur wage cuts. Taken together, the patterns associated with transient states and the vulnerable worker type align with the characteristics

Table 14: MEAN AND MEDIAN REAL WAGE CHANGES BY JOB TRANSITION TYPE

Transition	Mean (Median) Wage Change
JTJ (Type 5 = 0)	0.055 (0.032)
JTJ (Type 5 = 1)	0.018 (0.000)
JTJ (Type 6 = 0)	0.059 (0.039)
JTJ (Type 6 = 1)	0.032 (0.000)
JTJ (Type 7 = 0)	0.063 (0.042)
JTJ (Type 7 = 1)	0.025 (0.000)
JTJ (Types 5/6/7 = 0)	0.067 (0.047)
JTJ (Types 5/6/7 = 1)	0.033 (0.000)
Job Stayer	0.010 (0.000)
Job Changer via Unemployment	−0.010 (0.000)

Note to figure: This table shows the mean and median wage changes for job-to-job transitioners who were in the transient states indicated in parentheses.
Source: Authors' calculation

expected under a godfather shock.

Gertler et al. (2020) extend the model of Tjaden and Wellschmied (2014) to a business cycle context, positing that the share of JTJ moves accounted for by godfather shocks is higher in recession than expansion and that this explains the cyclical variation in the average wage change following a JTJ move relative the cyclicality exhibit by stayers and accessions from unemployment. To see the logic observe that if λ_d rises, holding all else equal, then the share of moves for which wages are a a draw from the full support of $F(w)$ rises and, thus, average wage changes associated with JTJ moves will be less positive. In other words, it may be the case that $F(w)$ is cyclically invariant while the wage change observed for the average JTJ mover fluctuates with the cycle simply due the varying proportion of moves that reveal preference v.s. are offers that can not be refused.

Our findings are not consistent with the Gertler et al. (2020) hypothesis. Specifically, our wage regressions reveal that the cyclicality of EE wage changes is heterogeneous with respect to a workers vulnerability at the time of transition. This stands in contrast to Tjaden and Wellschmied (2014) and Gertler et al. (2020), who posit that the distribution from which the godfather shock is drawn is acyclical. To accommodate our findings we build a new model of godfather shocks. A

companion paper by Ahn and Doniger (2026) can be one approach.

8 Conclusion

This paper investigates the effects of latent worker heterogeneity on wage changes among new hires. We uncover latent labor-market states and unobserved worker types by estimating a *multi-channel latent Markov model* using individual histories of labor force status and job switches under various circumstances from the Survey of Income and Program Participation. We identify highly transient latent labor-market states, typically experienced by workers with unstable employment who frequently switch jobs or become unemployed. Individuals who encounter transient states within a year or less are classified as the vulnerable type, with all others classified as the stable type. The vulnerable type—comprising 20 percent of the civilian population—does not experience wage gains when switching jobs, and their wages at new jobs exhibit significant excess cyclicality, in contrast to typical job switchers who experience wage gains and acyclical wage changes. Long tenure among vulnerable workers further reduces wage gains at new jobs and amplifies wage cyclicality, likely reflecting the destruction of job-specific human capital and greater mismatch. We interpret these empirical findings through the lens of a quantitative on-the-job search model with godfather shocks.

References

Abowd, John M., Francis Kramarz, and David N. Margolis, "High Wage Workers and High Wage Firms," *Econometrica*, March 1999, *67* (2), 251–333.

Ahn, Hie Joo, "The role of observed and unobserved heterogeneity in the duration of unemployment," *Journal of Applied Econometrics*, January 2023, *38* (1), 3–23.

_ **and Cynthia Doniger**, "Moving fast and moving slow. What does time-to-transition reveal about job opportunities?," *mimeograph*, 2026.

_ **and James D. Hamilton**, "Heterogeneity and Unemployment Dynamics," *Journal of Business & Economic Statistics*, 2020, *38* (3), 554–569.

_ **and** _ , "Measuring labor-force participation and the incidence and duration of unemployment," *Review of Economic Dynamics*, 2022, *44*, 1–32.

_ **and Jeremy B. Rudd**, "(Re-)Connecting Inflation and the Labor Market: A Tale of Two Curves," *Journal of Monetary Economics*, 2025, *153*, 103796.

_ **and Ling Shao**, "The Cyclicality of On-the-Job Search Effort," *The B.E. Journal of Macroeconomics*, 2021, *21* (1), 185–220.

_ **, Bart Hobijn, and Aysegul Sahin**, "The Dual U.S. Labor Market Uncovered," Finance and Economics Discussion Series 2023-031, Board of Governors of the Federal Reserve System (U.S.) May 2023.

Allman, Elizabeth S., Catherine Matias, and John A. Rhodes, "Identifiability of parameters in latent structure models with many observed variables," *Annals of Statistics*, 2009, *37* (6A), 3099–3132.

Bacci, S., S. Pandolfi, and F. Pennoni, "A comparison of some criteria for states selection in the latent Markov model for longitudinal data," *Advances in Data Analysis and Classification*, June 2014, *8* (2), 125–145.

Bacci, Silvia, Silvia Pandolfi, and Fulvia Pennoni, "A Comparison of Some Criteria for States Selection in the Latent Markov Model for Longitudinal Data," *Advances in Data Analysis and Classification*, 2014, *8* (2), 125–145.

Bartolucci, Francesco, Alessio Farcomeni, and Fulvia Pennoni, *Latent Markov Models for Longitudinal Data* 01 2013.

_ **, Silvia Bacci, and Fulvia Pennoni**, "Longitudinal Analysis of Self-Reported Health Status by Mixture Latent Auto-Regressive Models," *Journal of the Royal Statistical Society: Series C*

(Applied Statistics), 2014, *63* (2), 267–288.

Borovičková, Katarína and Claudia Macaluso, "Heterogeneous job ladders," *Journal of Monetary Economics*, 2025, *150*, 103711. Labor Market Dynamics: Drivers and Consequences.

Cairó, Isabel and Tomaz Cajner, "Human Capital and Unemployment Dynamics: Why More Educated Workers Enjoy Greater Employment Stability," *The Economic Journal*, 2018, *128*, 652–682.

Carrillo-Tudela, Carlos, "Job Search, Human Capital and Wage Inequality," Technical Report No. 6949, IZA Discussion Paper, Institute of Labor Economics 2012. Also available as CESifo Working Paper Series No. 3979 and ISER Working Paper Series No. 2012-23.

Doniger, Cynthia L., "Education and the Margins of Cyclical Adjustment in the Labor Market," 2023.

_ , "Wage dispersion with heterogeneous wage contracts," *Review of Economic Dynamics*, 2023, *51*, 138–160.

_ **and Desmond Toohey**, "These Caps Spilleth Over: Equilibrium Effects of Unemployment Insurance," Technical Report 2022-074, Federal Reserve Board, Finance and Economics Discussion Series November 2022. Available at SSRN.

Flinn, Christopher and Joseph Mullins, "Firms' Choices of Wage-Setting Protocols," Technical Report.

Fujita, Shigeru, "An empirical analysis of on-the-job search and job-to-job transitions," Working Paper 10–34, Federal Reserve Bank of Philadelphia 2010.

_ **, Giuseppe Moscarini, and Fabien Postel-Vinay**, "Measuring Employer-to-Employer Reallocation," NBER Working Papers 27525, National Bureau of Economic Research, Inc July 2020.

Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari, "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," *Review of Economic Studies*, 2020, *87* (4), 1876–1914.

Gregory, Victoria, Guido Menzio, and David Wiczer, "The alpha beta gamma of the labor market," *Journal of Monetary Economics*, 2025, *150*, 103695. Labor Market Dynamics: Drivers and Consequences.

Grigsby, John R., "Skill Heterogeneity and Aggregate Labor Market Dynamics," Working Paper 30052, National Bureau of Economic Research May 2022.

Hall, Robert E and Marianna Kudlyak, "Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics," Working Paper 25625, National Bureau

of Economic Research February 2019.

_ **and** _ , "The unemployed with jobs and without jobs," *Labour Economics*, 2022, *79*, 102244.

Hornstein, Andreas, Per Krusell, and Giovanni L. Violante, "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *American Economic Review*, December 2011, *101* (7), 2873–98.

Jolivet, Gregory, Fabien Postel-Vinay, and Jean-Marc Robin, "The empirical content of the job search model: Labor mobility and wage distributions in Europe and the US," *European Economic Review*, 2006, *50* (4), 877–907.

Juang, B. H. and Lawrence R. Rabiner, "Hidden Markov Models for Speech Recognition," *Technometrics*, 1991, *33* (3), 251–272.

Kudlyak, Marianna, "The cyclicality of the user cost of labor," *Journal of Monetary Economics*, 2014, *68* (C), 53–67.

_ , "How Cyclical Is the User Cost of Labor?," *Journal of Economic Perspectives*, 2024, *38* (2), 159–180.

_ **and Fabian Lange**, "Measuring Heterogeneity in Job Finding Rates Among the Nonemployed Using Labor Force Status Histories," *IZA Discussion Papers*, 2014, (No. 8663). Also available as FRB Richmond Working Paper No. 14-18 and on SSRN.

Low, Hamish, Costas Meghir, and Luigi Pistaferri, "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 2010, *100* (4), 1432–1467.

Morchio, Iacopo, "Work Histories and Lifetime Unemployment," *International Economic Review*, 2020, *61* (3), 1199–1235.

Mueller, Andreas I. and Johannes Spinnewijn, "The Nature of Long-Term Unemployment: Predictability, Heterogeneity and Selection," NBER Working Papers 30979, National Bureau of Economic Research, Inc February 2023.

Ozkan, Serdar, Jae Song, and Fatih Karahan, "Anatomy of Lifetime Earnings Inequality: Heterogeneity in Job-Ladder Risk versus Human Capital," *Journal of Political Economy Macroeconomics*, 2023, *1* (3), 506–550.

Postel-Vinay, Fabien and Jean-Marc Robin, "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, November 2002, *70* (6), 2295–2350.

_ **and** _ , "To match or not to match?: Optimal wage policy with endogenous worker search intensity," *Review of Economic Dynamics*, 2004, *7* (2), 297–330.

Robin, Jean-Marc, "On the Dynamics of Unemployment and Wage Distributions," *Econometrica*,

2011, *79* (5), 1327–1355.

Shibata, Ippei, "Labor Market Dynamics: A Hidden Markov Approach," IMF Working Papers 2019/282, International Monetary Fund December 2019.

Simmons, Michael, "Precautionary on-the-job search," SSRN Working Paper January 2024. Available at SSRN: https://ssrn.com/abstract=4685118.

Tanaka, Satoshi, Lawrence F. Warren, and David Wiczer, "Earnings Growth, Job Flows and Churn," *Journal of Monetary Economics*, 2023, *135*, 86–98.

Tjaden, Volker and Felix Wellschmied, "Quantifying the Contribution of Search to Wage Inequality," *American Economic Journal: Macroeconomics*, January 2014, *6* (1), 134–61.

Viterbi, Andrew J., "Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm," *IEEE Transactions on Information Theory*, 1967, *13* (2), 260–269.

A Additional Details about MLMM

A.1 Forward–Backward Algorithm for a Multi-Channel Latent Markov Model

Let $U_t \in [u_1, u_2, \cdots, u_k]$ be the latent state at time t and let the two emissions be Y_t^1 and Y_t^2. For the ease of communication, we suppress the notation for individual i for latent state variable U and observed emissions Y_t^1 and Y_t^2. The parameters are:

- Initial state probabilities: $\pi_1 = P(U_1 = u)$
- Transition probabilities: $\Pi = P(U_t = u | U_{t-1} = u') \quad t = 2, 3, ..., T, \quad u', u \in [u_1, u_2, \cdots, u_k]$.
- Emission densities: ϕ_u^1 and ϕ_u^2

We assume conditional independence of channels given $U_t = u$:

$$p\left(Y_t^1, Y_t^2 \mid U_t = u\right) = \phi_u^1 \phi_u^2.$$

Forward Recursion

Define the forward probability

$$\alpha_t(u) = P\left(Y_{1:t}^1, Y_{1:t}^2, U_t = u \mid \theta\right),$$

where $Y_{1:t}^1$ and $Y_{1:t}^2$ are the observed sequences of two emissions from the first period to time t and the vector of model parameters is denoted as θ.

The forward component is initialized ($t = 1$) as follows:

$$\alpha_1(u) = \pi_1(u)\, \phi_u^1(Y_t^1) \phi_u^2(Y_t^2), \quad u \in [u_1, u_2, \cdots, u_k].$$

The recursion is computed as follows ($t = 2, \ldots, T$):

$$\alpha_t(u) = \left[\sum_{u' \in [u_1, \cdots, u_k]} \alpha_{t-1}(u') \pi_{u'u} \right] \phi_u^1(Y_t^1) \phi_u^2(Y_t^2),$$

where $\pi_{u'u}$ is a Markovian transition probability between u' and u.

The likelihood is computed as follows:

$$\mathscr{L}(\theta) = \sum_{u' \in [u_1, \cdots, u_k]} \alpha_T(u).$$

Backward Recursion

Define the backward probability:

$$\beta_t(u) = P\left(Y_{t+1:T}^1, Y_{t+1:T}^2 \mid U_t = u, \theta\right),$$

where $Y_{t+1:T}^1$ and $Y_{t+1:T}^2$ are the observed sequences of two emissions between $t+1$ and T, and the vector of model parameters is denoted as θ.

For the initialization of the backward component ($t = T$), we set

$$\beta_T(u) = 1, \quad u \in [u_1, \cdots, u_k].$$

The backward recursion is computed as follows for $t = T-1, \ldots, 1$:

$$\beta_t(u') = \sum_{u \in [u_1, \cdots, u_k]} \pi_{u'u} \phi_u^1\left(Y_{t+1}^1\right) \phi_u^2\left(Y_{t+1}^2\right) \beta_{t+1}(u).$$

State and Transition Posteriors

With the forward and backward components, we compute the state and transition posteriors.

1. State posterior:

$$\gamma_t(u) = P(U_t = u \mid Y^1_{1:T}, Y^2_{1:T}, \theta) = \frac{\alpha_t(u)\beta_t(u)}{\mathscr{L}(\theta)}.$$

2. Transition posterior:

$$\xi_t(u', u) = P(U_t = u', U_{t+1} = u \mid Y^1_{1:T}, Y^2_{1:T}, \theta) = \frac{\alpha_t(u')a_{u'u}\phi^1_u\left(Y^1_{t+1}\right)\phi^2_u\left(Y^2_{t+1}\right)\beta_{t+1}(u)}{\mathscr{L}(\theta)}.$$

Scaling for Numerical Stability

To avoid underflow, scale $\alpha_t(u)$ and $\beta_t(u)$ at each t by a factor c_t, and compute:

$$\ell(\theta) = -\sum_{t=1}^{T} \log c_t$$

where $\ell(\theta)$ is the log-likelihood value.

A.2 Model Selection: Number of Latent States

The two statistics are computed from the following:

$$\begin{aligned} \text{AIC} &= -2\hat{\ell}(\theta) + 2 \times \#\text{ of parameters} \\ \text{BIC} &= -2\hat{\ell}(\theta) + 2log(N) \times \#\text{ of parameters} \end{aligned}$$

where $\hat{\ell}(\theta)$ is the estimated maximum of the log-likelihood value given the parameter estimates $\hat{\theta}$ and N is the number of observations. The number of latent states that yields the minimum value of AIC and BIC is chosen. Both model selection criteria support seven latent states. This model represents the best compromise between goodness-of-fit and model complexibility (the associated parameter uncertainty). Bartolucci et al. (2013) recommends to use BIC if the two criteria conflict, but also suggest to consider other criteria such as those taking into account the quality of the classification (see, for example, Bartolucci, Bacci and Pennoni, 2014; Bacci,

Pandolfi and Pennoni, 2014b).

A.3 Decoding

There are two types of decoding: local decoding and global decoding. Local decoding captures the best predicted latent state of individual i at time t given information set through t. Global decoding captures the best-predicted latent state of individual i at time t after seeing the entire sequence of responses through the end of the sample period T. Local decoding is analogous to one-sided inference and global decoding corresponds to two-sided inference in Kalman filter. With the model estimates, we uncover both locally-decoded and globally-decoded histories of latent states for each individual over a span of 64 months.[20]

(Local Decoding) The EM algorithm directly provides the estimated posterior probabilities of U^t, denoted by

$$P\left(U^t = u \mid \tilde{\mathbf{Y}} = \tilde{y}\right),$$

for $t = 1,\ldots,T$, $u = [u_1, u_2, \ldots, u_k]$, and for every realized response $(\tilde{y})$ observed at least once. These probabilities are maximized to obtain a prediction of the latent state of each subject at each time t. Note that this type of decoding minimizes the classification error at each time t, but may yield sub–optimal predictions of $U^1,\ldots,U^T$.

(Global Decoding) The most likely posterior sequence of states after seeing the data through T is obtained through the so-called *global decoding*. The global decoding is based on an adaptation of the Viterbi (1967) algorithm; see also Juang and Rabiner (1991). The algorithm proceeds through a forward–backward recursion within the maximum-likelihood estimation via the EM algorithm, and is performed for long sequences of data.[21]

[20] Details on the local and global decoding are largely from Bartolucci et al. (2013).
[21] Bartolucci et al. (2013) discuss the details.