Journal of Public Economics 197 (2021) 104398

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

Projecting unemployment durations: A factor-flows simulation approach with application to the COVID-19 recession ☆

Gabriel Chodorow-Reich [a,*], John Coglianese [b]

[a] Harvard, United States
[b] Board of Governors of the Federal Reserve System, United States

ARTICLE INFO

Article history:
Received 22 July 2020
Revised 14 December 2020
Accepted 23 February 2021
Available online 5 March 2021

Keywords:
Unemployment duration
Labor market simulation
COVID recession

ABSTRACT

We propose a three-step factor-flows simulation-based approach to forecast the duration distribution of unemployment. Step 1: estimate individual transition hazards across employment, temporary layoff, permanent layoff, quitter, entrant, and out of the labor force, with each hazard depending on an aggregate component as well as an individual's labor force history. Step 2: relate the aggregate components to the overall unemployment rate using a factor model. Step 3: combine the individual duration dependence, factor structure, and an auxiliary forecast of the unemployment rate to simulate a panel of individual labor force histories. Applying our approach to the November Blue Chip forecast of the COVID-19 recession, we project that 750,000 workers laid off in April 2020 remain unemployed eight months later. Total long-term unemployment rises thereafter and eventually reaches 4.2 million individuals unemployed for more than 26 weeks and 1.4 million individuals unemployed for more than 46 weeks. Long-term unemployment rises even more in a more pessimistic recovery scenario, but remains below the level in the Great Recession due to a high amount of labor market churn.

1. Introduction

Long-term unemployment poses unique challenges to society. Consistent with standard incomplete insurance models, as jobless spells lengthen, individuals exhaust their savings (Bewley, 1980; Ganong and Noel, 2019). The possibility of a deterioration of skills threatens productivity even after the economy returns to full employment. And the duration structure of unemployment matters directly to policy, as every recession in the United States in the past 50 years has included an extension of the number of weeks individuals may receive government-provided unemployment insurance (UI) benefits. The question of how many weeks to extend benefits depends in part on how many individuals will reach different unemployment durations.

Projecting the duration structure of unemployment also poses challenges. Even conditional on a forecast of the overall unemployment rate, the duration structure can vary widely depending on the magnitude of gross flows across labor force states. For example, ignoring for the moment transitions in and out of the labor force, an unemployment rate of 9% could arise from a monthly separation rate into unemployment of 5% and job-finding rate out of unemployment of 50%, or a separation rate of 2% and job-finding rate of 20%. A job-finding rate of 50% implies an ergodic distribution in which fewer than 2% of unemployed individuals have been unemployed for more than six months, while a job-finding rate of 20% implies more than one-quarter of unemployed individuals have a duration greater than six months. Furthermore, the duration distribution also depends on the amount of individual duration dependence in job finding rates.

In this paper, we introduce a factor-flows simulation-based approach that addresses these difficulties. The approach has three steps. In the first step, we estimate individual transition hazards across six labor force states: employment, unemployment on temporary layoff, permanently laid-off, quitters, entrants, and out of the labor force. The estimated hazards depend on an individual's labor force history, consistent with recent work emphasizing the path dependence in hazard rates (Jarosch, 2015; Yagan, 2019; Kudlyak and Lange, 2018; Coglianese, 2018; Hall and Kudlyak, 2019), as well as on an aggregate component. In the second step, we specify a factor model for the aggregate components of the hazard rates. In practice, the first principal component explains a large share of the variation across components and has a correlation of 0.97 with the unemployment rate. We therefore treat the overall unemployment rate as an observed factor. In the third step, we

☆ We are grateful to Johannes Spinnewijn and two anonymous referees for their comments. The views expressed herein are those of the authors and not necessarily those of the Board of Governors of the Federal Reserve System.
* Corresponding author.
E-mail address: chodorowreich@fas.harvard.edu (G. Chodorow-Reich).

https://doi.org/10.1016/j.jpubeco.2021.104398
0047-2727/

obtain a forecast for the overall unemployment rate from an auxiliary source and use the estimated factor model and individual duration dependence in transition hazard rates to simulate a panel of individual labor force histories consistent with the auxiliary forecast of the overall unemployment rate, the implied aggregate components of the transition hazards, and the historical pattern of duration dependence.

We apply our approach to the COVID-19 recession. The COVID-19 pandemic impacted the U.S. labor market in an unprecedented fashion. From a low of 3.5% in February 2020, the official unemployment rate rose to a post-war high of 14.7% in April.[1] The largest prior two-month increase in the unemployment rate was 1.5 p.p. in 1975. The 7.8 p.p. decline in the unemployment rate over the subsequent six months also marks a historical record. Nonetheless, the sheer scale of the increase in unemployment raises the possibility of a large cohort of individuals that remain unemployed for a substantial period of time.

Fig. 1 shows the duration distribution in our baseline simulation, separately for those on layoff (temporary layoff U^t and permanent layoff U^p) and other unemployed individuals (quitters U^q and entrants U^e). The average transition hazards and unemployment rate through October 2020 come from the Current Population Survey (CPS), the source for the official unemployment rate. Despite the historically rapid recovery following the dramatic increase in unemployment in March and April, a shrinking yet by-historical-standards large stock of individuals have remained without work for more than six months, as shown in the figure. After October 2020, the overall unemployment rate follows the November 2020 Blue Chip Consensus Forecast, which averages the unemployment rate forecasts from more than 50 leading business economists. According to the simulation, total long-term unemployment eventually reaches 4.2 million individuals unemployed for more than 26 weeks and 1.4 million individuals unemployed for more than 46 weeks.[2]

One crucial parameter governing the duration distribution is the re-employment hazard of workers on temporary layoff. More than 90% of the individuals transitioning from employment to unemployment in April 2020 reported being on temporary layoff; in October 2020 there were still 3 million such individuals, compared to around 0.5 million pre-COVID. According to the CPS classification scheme, these individuals were available to work in the survey reference week and had either received a recall date from their employer or had been given an indication that they would be recalled to work within the next six months. Over the period 1994 to 2019, roughly 50% of such individuals returned to employment the following month, more than double the re-employment hazard for unemployed individuals not on temporary layoff. However, the re-employment hazard has historically declined especially precipitously for this group — individuals in their first month of temporary layoff have been re-employed at a rate of 58%, while those on temporary layoff for two months have a re-employment rate of 43%. Our simulation incorporates this duration dependence, increasing the amount of long-term unemployment commensurately.

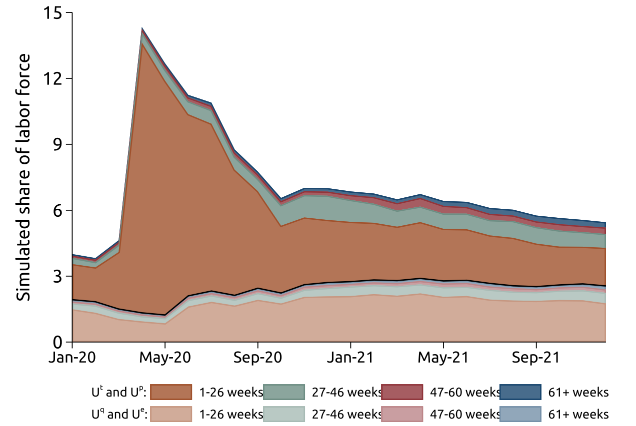

Fig. 1. Projected Unemployment Duration Distribution. *Notes:* The figure plots the simulated share of the labor force laid off (U^t and U^p) and other unemployed (U^q and U^e), by the number of weeks in the past two years the individual has been unemployed. The overall path of unemployment follows the November Blue Chip Consensus.

The implementation just described contains a number of judgment calls. These include the functional form of the hazard equations, the period over which to estimate them, and the specification of the factor model, among others. We assess our choices using two out-of-sample exercises. The first uses a hold-out sample to compare our specification of transition hazards to several alternatives. The second repeats the forecasting exercise in the 2001 recession, the Great Recession, and June-October 2020 and compares the simulation unemployment rate and duration distribution to the CPS.

We also simulate two alternative scenarios to assess the sensitivity of the forecast for the COVID period. The *pessimistic scenario* follows the average of the ten highest Blue Chip unemployment rate forecasts and envisages an unemployment rate 0.8 p.p. higher than the Consensus throughout 2021. The *optimistic scenario* follows the average of the ten lowest Blue Chip forecasts and envisages an unemployment rate 0.8 p.p. lower throughout 2021. Strikingly, the amount of long term unemployment does not differ dramatically across these scenarios. This result reflects the high degree of labor market churn across simulations, irrespective of the path of unemployment.

Our paper complements other research studying the labor market in the early months of the COVID recession (Bartik et al., 2020; Cajner et al., 2020). Whereas this literature has examined trends to date, our focus is on future implications, although we necessarily also cover new ground on the labor market dynamics in the early months of the recession. Similar to Barrero et al. (2020), our baseline simulation contains a substantial amount of labor market reshuffling, despite using a very different methodology. Our methodology also builds on the aforementioned literature emphasizing duration dependence in labor market transitions, although we do not take a stand on the source of this dependence (Baker, 1992; Krueger et al., 2014; Kroft et al., 2013; Kroft et al., 2016; Jarosch and Pilossoph, 2018; Mueller et al., 2018; Ahn and Hamilton, 2020a). We are not aware of previous research empirically associating hazard rates across labor market states with a factor structure, although the basic search-and-matching model has this feature.

Section 2 describes our methodology in detail and presents intermediate results including the underlying transition hazards. Section 3 contains our main results. Section 4 discusses implications for UI policy.

[1] According to the Bureau of Labor Statistics, in April survey respondents representing 7.5 million individuals who reported being temporarily absent from their jobs were possibly mis-classified as employed instead of unemployed (https://www.bls.gov/cps/employment-situation-covid19-faq-may-2020.pdf). Table A.1 reports the number potentially mis-classified by month, those receiving pay, and the average transition rate into employed at work. See https://www.brookings.edu/blog/up-front/2020/06/30/who-are-the-potentially-misclassified-in-the-employment-report/ for further discussion.

[2] Unemployment durations in this paragraph and throughout the paper refer to the total number of weeks an individual has been unemployed over the previous two years. This definition differs from the official BLS definition of consecutive weeks of unemployment. The numbers are based on a total civilian, non-institutional population 16 and over of 260 million.

2. Methodology

Our methodology builds on the unemployment duration simulations in Chodorow-Reich and Coglianese (2019) by making several important improvements and adjustments tailored to the unique aspects of the COVID-19 recession, including allowing for multiple categories of unemployment and the explicit linking of the hazard rate factor model to unemployment rate loadings. We therefore provide a comprehensive discussion here, recognizing the foundations laid in that earlier work.

2.1. Hazard estimation

The first step of our methodology uses longitudinally-matched monthly data from the Current Population Survey (CPS) to estimate hazard rates for transitions across labor market states. Households drawn into the CPS sample are interviewed up to eight times over a sixteen month period about the labor force status of their members, with four consecutive months in sample initially followed by eight months out of sample and then four final consecutive months in sample, sometimes referred to as the 4-8-4 rotation group design. We restrict the sample to respondents who complete all eight interviews and re-weight to account for non-random attrition in responses.[3]

Our methodology can accommodate an arbitrarily flexible partition of labor force states, including by demographic group or occupation. For our application to the COVID recession, we model six: employed (E), unemployed-temporary layoff (U^t), unemployed-permanent layoff (U^p), unemployed-quit (U^q), unemployed-entrant (U^e), and not in the labor force (N). Fig. 2 illustrates why we split unemployment into these sub-categories. The historical hazard rates into employment and the incidences of these different categories of unemployment at the start of the COVID recession differ markedly. Remarkably, the figure reveals the increase in hiring during the labor market recovery to be an entirely compositional phenomenon; re-employment hazards within both temporary and permanent layoff remained below their historical average, but the high share of temporary layoffs and the historically higher re-employment hazard in this group pulled up the overall hiring rate. Ignoring this dimension of heterogeneity would cause our simulation to mis-characterize the amount of labor market churn.

We model each of the thirty-six hazard rates governing the transitions across the six states $s \in \{E, U^t, U^p, U^q, U^e, N\}$ as a function of an individual's labor market history and aggregate economy wide trends. Let $\gamma^s_{i,t}$ denote an indicator for individual i in month t being in state s, $\gamma^{s' \to s''}_{i,t}$ an indicator for transition from state s' in period $t-1$ to state s'' in period t, and $\zeta^s_{i,t-3}$ the self-reported duration of unemployment in weeks (zero if not unemployed) in period $t-3$ to capture labor force status while the individual had rotated out of the CPS sample. We estimate the OLS regression:

$$\gamma^{s' \to s''}_{i,t} = \sum_{\ell \in \{2,3,12,13,14,15\}} \sum_{s \in \{E,U^t,U^p,U^q,U^e,N\}} \phi^{s' \to s''}_{s,\ell} \gamma^s_{i,t-\ell} + \psi^{s' \to s''} \zeta^s_{i,t-3} + \bar{\delta}^{s' \to s''}_t + \epsilon^{s' \to s''}_{i,t}. \tag{1}$$

The terms $\phi^{s' \to s''}_{s,\ell} \gamma^s_{i,t-\ell}$ and $\psi^{s' \to s''} \zeta^s_{i,t-3}$ capture the history-dependence of transition hazards, using the available information about an individual's labor market history in the CPS.[4] The term $\bar{\delta}^{s' \to s''}_t$ is a month fixed effect that reflects the state of the aggregate labor market.

We estimate Eq. (1) over the period 1994-2020 and report the coefficients in Fig. A.3 of the appendix. For both temporary and permanent layoff, having been employed in the prior two months predicts a higher likelihood of re-employment. Differences also exist; for example, conditional on recent employment history, employment a year ago raises the likelihood of re-employment from temporary layoff but not from permanent layoff, and a history of temporary rather than permanent layoff increases the re-employment hazard for those currently on permanent but not for those currently on temporary layoff. History also matters to separation probabilities. Individuals with previous spells of unemployment have a higher separation hazard, and the excess probability concentrates into a return to their previous type (U^t or U^p) of unemployment.

2.2. Factor model

In the second step, we specify a factor structure for the aggregate components of the transition hazards:

$$\bar{\delta}^{s' \to s''}_t = \alpha^{s' \to s''} + \beta^{s' \to s''} F_t + v^{s' \to s''}_t, \tag{2}$$

where the factor loading $\beta^{s' \to s''}$ encodes the sensitivity of the aggregate component of the transition rate from s' to s'' to the common factor F_t. The factor structure has the key advantage of fixing the ratio of changes in the hazard rates, e.g. when the $E \to U^t$ hazard rises by 1 p.p., the $E \to N$ hazard will rise by $\frac{\beta^{E \to N}}{\beta^{E \to U^t}}$ p.p. In this way, it reduces the dimensionality of the aggregate component of the simulation from 36 separate transition hazards to a single factor.

For this dimensionality reduction to work in practice, a single factor must adequately capture the variation in the transition hazards and this factor must be linked to the auxiliary unemployment rate forecast. We start by performing principal components analysis on the 12-month moving averages of the 36 aggregate components of the transition hazards (to remove sampling volatility). The first principal component explains 29% of the variance of these series, more than double the share explained by the second principal component. Therefore, although even the 12-month moving averages contain substantial idiosyncratic volatility, a single factor explains the transition hazards well.

Fig. 3 plots the first principal component against the overall unemployment rate. The two series co-move extremely closely, with a correlation coefficient of 0.97. Based on this evidence, we directly equate the factor F_t with the overall unemployment rate and estimate the loadings $\beta^{s' \to s''}$ from regressions of the aggregate components of the transition hazards on the unemployment rate. Conveniently, the auxiliary forecast of the unemployment rate

[3] Fig. A.1 shows that respondents who complete all eight interviews have a lower average unemployment rate than other respondents, and that this bias mostly reflects the lower unemployment rate among any individual interviewed in consecutive months. The BLS produces research series for gross flows across employment, unemployment, and out of the labor force that correct for non-random attrition (https://www.bls.gov/cps/cps_flows.htm). We re-weight the individuals in our sample to match these corrected gross flows, applying the same scaling factor to all flows into and out of unemployment states. In this way, our sample of respondents completing their eighth interview matches overall unemployment dynamics. See Krueger et al. (2017) and Ahn and Hamilton (2020b) for further discussion of sample attrition in the CPS. An additional complication arises in the early months of the COVID recession, as response rates fell dramatically for cohorts entering the CPS sample for the first time, likely reflecting the temporary discontinuance of in-person interviews for these households. Fig. A.2 reports the response rates and unemployment rates by month-in-sample for 2020. Despite the much lower response rates of cohorts entering the sample in March-August 2020, the unemployment rate in these cohorts looks quite similar to the overall unemployment rate, so we make no additional adjustment for these months. Throughout the paper, we use the CPS longitudinal links created by Drew et al. (2014) and data from Flood et al. (2020) as well as from the basic monthly CPS files provided by the Census Bureau.

[4] To avoid overfitting, we include $\zeta^s_{i,t-3}$ only in the regressions for which $s' \in \{U^t, U^p, U^q, U^e\}$. We have confirmed that this approach produces better out-of-sample fit than alternatives, as described in section 2.3.

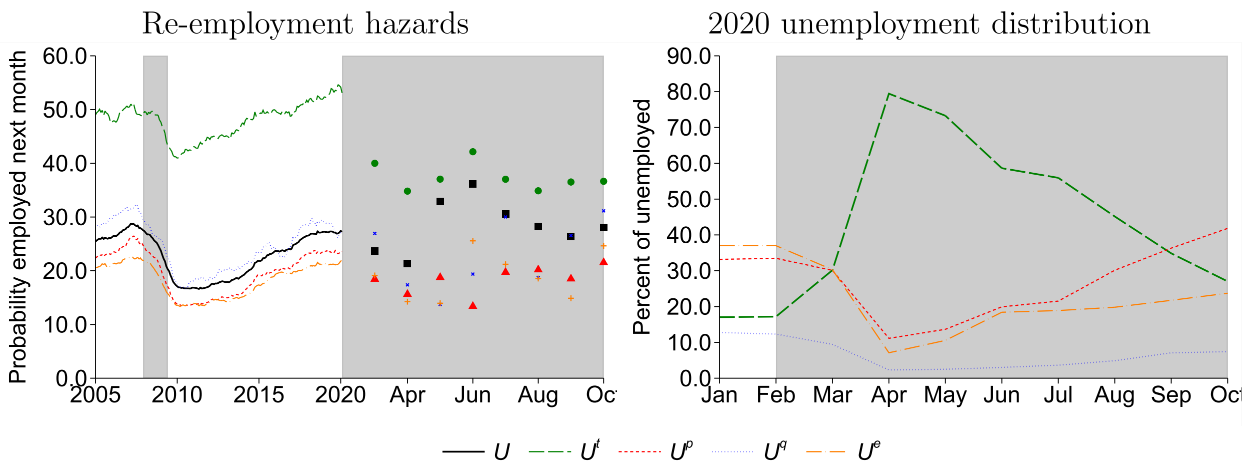

Fig. 2. Re-employment Heterogeneity. *Notes:* The left panel plots the re-employment probabilities from unemployment overall (U) and the sub-categories unemployed-temporary layoff (U^t), unemployed-permanent layoff (U^p), unemployed-quit (U^q), and unemployed-entrant (U^e) as twelve month moving-averages through February 2020 and the monthly values thereafter. The right panel plots the distribution of unemployment by status in 2020.

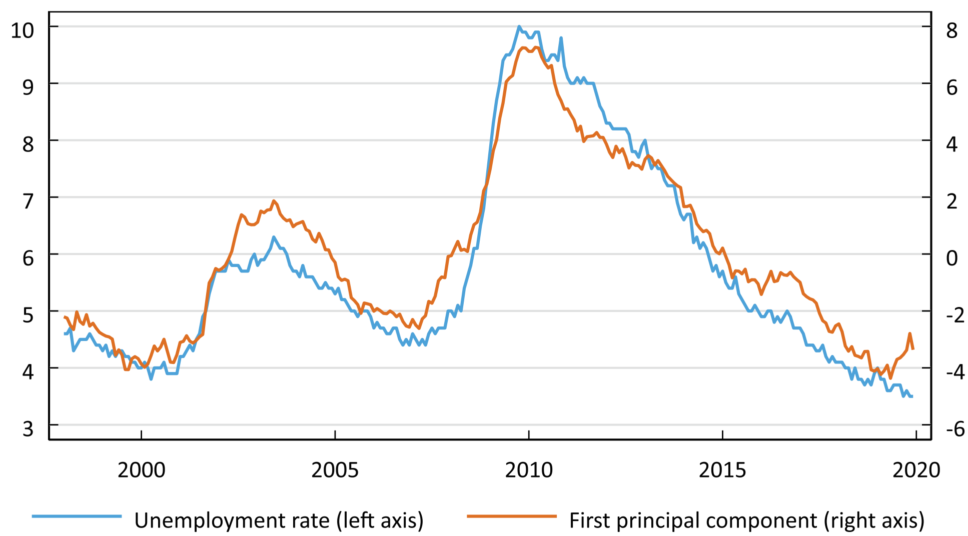

Fig. 3. Single Factors for Aggregate Transition Hazard Components. *Notes:* The figure plots the monthly unemployment rate (left-axis) with the first principal component of the 12-month moving averages of the 36 aggregate components of the transition hazards (right-axis, scale inverted for comparison).

then directly determines the aggregate components of the transition hazards. We report the factor loadings in table A.2.

2.3. Model validation

The factor-flow approach may incorrectly forecast unemployment duration for three main reasons. First, it relies on the accuracy of the forecast of the forcing variable F_t. This source of error is intrinsic to any forecasting exercise. Second, there may be model mis-specification, for example if the coefficients $\{\phi_{s,\ell}^{s'\to s''}, \psi^{s'\to s''}, \beta^{s'\to s''}\}$ are actually time-varying.[5] Third, the structural residual $\nu_t^{s'\to s''}$ is a non-idiosyncratic shock to flows in each period that will cause the simulation to depart from the data even if Eqs. (1) and (2) are correctly specified and the coefficients consistently estimated.

We conduct a cross-validation exercise to compare our specification to possible alternatives. We divide CPS respondents into four equal-sized groups that are evenly distributed across months. For each 1/4 group, we estimate Eqs. (1) and (2) on the other 3/4 of the sample and use these estimates to construct predicted probabilities $\hat{\gamma}_{i,t}^{s'\to s''}$ for transitions in the 1/4 group. Our cross-validation metric is the average negative log likelihood of the observed transitions $CV = -\log(\hat{\gamma}_{i,t}^{s'\to s''})$ in the sample, where lower values indicate better out-of-sample fit. We repeat this procedure for several alternative specifications, including with additional variables, interactions, substituting Eq. (2) into Eq. (1) and estimating transition probabilities in a single step, and estimating the individual heterogeneity component in different sub-periods, and report CV for each in Appendix Table A.3, both for the full sample and for individual recessionary periods.

This exercise demonstrates the flexibility of our approach, but also the range of choices required to implement it. Cross-validation offers a disciplined way to choose among alternatives. Our baseline specification obtains better out-of-sample fit than

[5] A recent literature views separation rates as causally affected by an unemployment spell (Jarosch, 2015; Hall and Kudlyak, 2019). If instead history dependence largely reflects selection of who becomes unemployed and the selection mechanism is muted in the COVID recession, then the historical values of the history dependence coefficients governing separation rates will be too large and our simulations will overstate the amount of history dependence and hence the amount of long-term unemployment.

the alternatives we examine, both for the sample as a whole and during the 2001 and 2008-09 recessions and May-July 2020 period.

2.4. Simulation

Having estimated $\phi_{\ell}^{s'\to s''}, \psi^{s'\to s''}$, and $\beta^{s'\to s''}$ from the CPS over the 1994–2020 period, we start the simulation with three burn-in phases to achieve an unemployment duration distribution that is approximately at steady-state. In the first phase, we initialize 100,000 individuals with randomly drawn initial states, using the ergodic distribution of $s_{i,t}$. In the second phase, we simulate an additional 13 months for each individual, drawing $s_{i,t}$ from a distribution conditional on $s_{i,t-1}$, where we use the national average transition probabilities over 1994-2020. Once we have 14 periods of data, we construct individual-specific transition probabilities. For each month and individual, we generate six probabilities for $s_{i,t}$ conditional on $s_{i,t-1}$ by applying the estimated duration dependence coefficients $\phi_{\ell}^{s'\to s''}, \psi^{s'\to s''}$ to the individual's simulated labor market history and the estimated aggregate loadings $\beta^{s'\to s''}$ to the target unemployment rate. We simulate labor market histories in this way for 24 months to complete the burn-in.

Starting from the simulated ergodic distribution, we use the actual estimates of the aggregate components $\bar{\delta}_t^{s'\to s''}$ for January 2016 through October 2020 to simulate the initial conditions and months of the COVID recession. For this period, we constrain the average labor market transition rates in the simulation to exactly match the observed rates from the CPS.[6] After October 2020, we set the factor so that the simulated unemployment rate matches the November 2020 Blue Chip Consensus forecast. The Blue Chip Economic Indicators is a monthly survey of more than 50 leading business economists conducted over two days at the start of each month. The Blue Chip Consensus is the average across all Blue Chip forecasters and features an unemployment rate of 7.6% in 2020Q4 and 6.2% in 2021Q4.

We also simulate alternative scenarios. The *pessimistic scenario* follows the average of the ten highest Blue Chip unemployment rate forecasts and envisages an unemployment rate 0.8 p.p. higher than the Consensus in 2020Q4 and throughout 2021. The *optimistic scenario* follows the average of the ten lowest Blue Chip forecasts and envisages an unemployment rate 0.8 p.p. lower in 2020Q4 and throughout 2021.

3. Results

In this section we discuss several noteworthy aspects of our application to the COVID-19 recession.

3.1. Prevalence of long-term unemployment

For each unemployed individual, we define months unemployed as the total amount of time that individual has spent in unemployment in the previous two years. This definition differs from the official BLS definition of consecutive months of unemployment. It has the advantage of mitigating the problem of persons spuriously reporting that they moved between unemployment and out of the labor force, which would also contaminate our simulation since it uses actual labor force histories as input data (see Ahn and Hamilton (2020b) for further discussion). It also better accords with policy questions such as unemployment insurance (UI) exhaustion, which depends on total weeks of benefit receipt even when interrupted by an employment spell.

In the baseline simulation, the number of individuals unemployed for longer than 26 weeks peaks in December 2020 at 4.2 million. Very long-term unemployment peaks about a year later in early 2022 with 1.4 million individuals unemployed for more than 46 weeks and 900,000 unemployed for longer than 60 weeks. Although high, these magnitudes fall below the corresponding peak levels in the Great Recession, despite a higher overall unemployment rate peak in the COVID period. This difference reflects the different conditions at the start of each episode.

3.2. Churn

Fig. 4a shows actual and projected gross hires and separations in each month. These flows remain elevated relative to their pre-recession levels. This pattern reflects two complementary forces. First, the high number of unemployed on temporary recall generates relatively high re-employment hazards, especially early in the recovery. Matching the path of the unemployment rate then requires that separations also remain high. Notably, initial claims for unemployment insurance at the end of November remained above their peak in the Great Recession. Despite not targeting UI claims data, our simulation matches the implied separation rate closely. Second, as an increasing number of individuals experience unemployment, the model perceives their attachment to their employer to decline. This deterioration reflects the historical tendency for individuals with recent non-employment to separate from their employer more frequently than those with long spells of employment. In the specific circumstance of the COVID recession, it could reflect workers in socially interactive jobs most vulnerable to health-related fluctuations in demand. The high amount of churn reduces the prevalence of long-term unemployment at a given level of the unemployment rate, even in our expansive definition that allows for temporary periods of re-employment.

Fig. 4b provides another perspective on the amount of churn by plotting the share of unemployed in each month that first entered unemployment in the April 2020 spike. In the actual data, about 11 million individuals transitioned from employment to unemployment in April 2020, with more than 90% of these individuals on temporary layoff. Among the April 2020 temporary layoff cohort, 40% were re-employed in May and 58% were re-employed in June. As a result, those entering temporary layoff in April 2020 accounted for less than one-third of total unemployment in June. Our simulation tracks these flows closely. On the other hand, those remaining unemployed face lower re-employment hazards, either because they have transitioned to permanent layoff or out of the labor force or because of the negative duration dependence among the temporary unemployed, and even those re-employed have higher separation hazards due to the history dependence. These characteristics explain why the April 2020 cohort remains a non-trivial share of total unemployment even in 2021.

We can also ask what the projected churn implies for reallocation dynamics. Overall, two-thirds of the April 2020 temporary layoff cohort eventually transition back to employment directly from temporary unemployment. Associating these individuals with recalls yields a share of the April 2020 cohort returning to their previous employer of two-thirds, in the range of the forecast made

[6] Specifically, for an individual with previous-month state s', let $\Pr(s'\to s'',i)$ denote the predicted probability i transitions to state s'' next month using Eqs. (1) and (2). We rescale this as $\tilde{\Pr}(s'\to s'',i) = \Pr(s'\to s'',i)\cdot\frac{\bar{\gamma}_t^{s'\to s''}}{\sum_{j\in s'}\Pr(s'\to s'',j)}$, where $\sum_{j\in s'}\Pr(s'\to s'',j)$ is the sum of predicted probabilities across all simulated individuals who were in state s' last month and $\bar{\gamma}_t^{s'\to s''}$ is the average $s'\to s''$ observed in the CPS in month t.

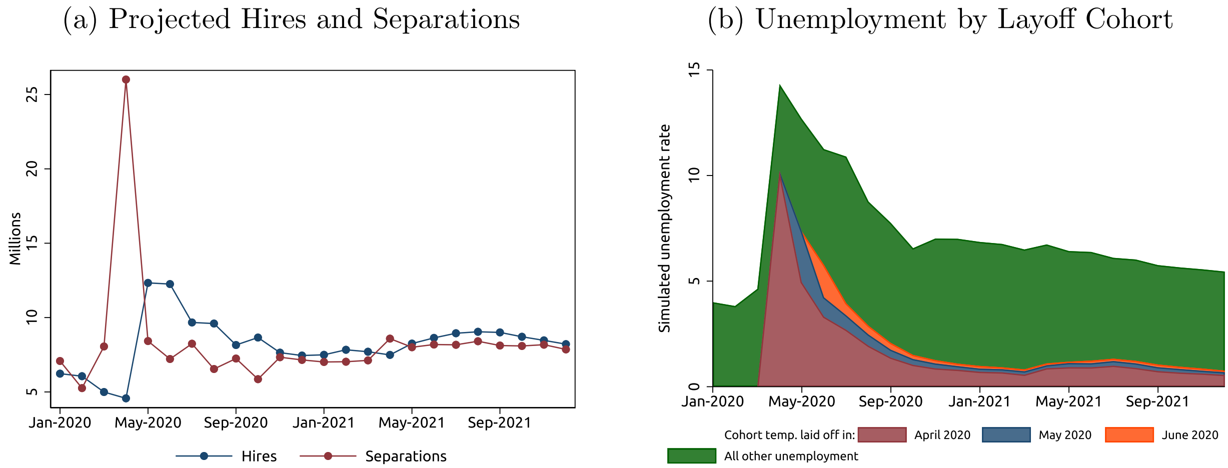

Fig. 4. Churn. *Notes:* The left figure plots the number of hires from non-employment and separations to non-employment in our baseline simulation. The right figure plots the portion of unemployment attributable to different layoff cohorts. Cohorts are defined based on $E \to U^t$ flows in each month, but all subsequent months of unemployment are counted as attributable to that cohort, including months of U^p, U^q, and U^e.

in Barrero et al. (2020) despite our using a very different methodology.[7]

3.3. Alternative scenarios

The Blue Chip Consensus forecast and our simulation involve a historically rapid labor market recovery. The left panel of Fig. 5 shows the duration distribution in a more pessimistic scenario in which the unemployment rate follows the average of the 10 highest Blue Chip forecasts and exceeds 8% in 2020Q4 before falling to 7.1% in 2021Q4. Perhaps surprisingly, despite the slower labor market recovery, long-term unemployment does not rise that much relative to the baseline scenario, with the number unemployed more than 6 months peaking at 5 million.

The right panel of Fig. 5 shows the duration distribution in an optimistic scenario in which the unemployment rate follows the average of the 10 lowest Blue Chip forecasts and falls to 7.2% in 2020Q4 and 5.3% in 2021Q4. This scenario contains less long-term unemployment, with the number unemployed for more than 6 months peaking at 4 million, but again does not differ dramatically from the baseline or even pessimistic scenario. The explanation lies in the assumed churn, which does not differ too much across all three scenarios. As a result, while the pessimistic scenario contains a lower hiring rate and higher separation rate than the baseline, and the optimistic scenario a higher hiring rate and lower separation rate, these differences have a limited impact on the amount of long-term unemployment. By contrast, if the labor market becomes more sclerotic than in our simulations, long-term unemployment could rise much more.

3.4. Out-of-sample fit

As another validation exercise, we examine how our method performs in pseudo out-of-sample exercises. For the 2001 recession, Great Recession, and June–October 2020, we simulate the ergodic distribution and use the observed transitions between states to simulate the periods leading up to the unemployment peak. Following the peak, we use the observed unemployment rate as the factor F_t to simulate the distribution as if it were a forecast.

Appendix Fig. A.5 shows the simulated unemployment duration distribution, simulated unemployment rate, the actual unemployment rate, and an "apples-to-apples" comparison of simulated and actual duration using the number of months unemployed for unemployed individuals in rotation group 8 computed the same in the simulation as in the data. In both the COVID episode and the 2001 recession, the simulated unemployment rate and the duration distributions closely track their actual values. In the Great Recession the simulated unemployment rate continues to rise after the actual unemployment rate peaks in October 2009. This discrepancy appears to reflect an unusual number of CPS respondents identifying as out of the labor force during this period; in fact, Ahn and Hamilton (2020b) argue that correcting for misreporting in the CPS produces less rapid declines in participation and in unemployment during this period than found in the official series, and hence closer to our simulated series.

4. Implications

We conclude by highlighting the implications for unemployment insurance (UI). In normal periods, individuals may claim UI benefits from their state insurance programs, typically for up to 26 weeks. The joint federal-state Extended Benefits (EB) program provides up to an additional 20 weeks of benefits in states with high unemployment. Finally, in every recession since 1950, the federal government has enacted temporary recipiency tiers that allow individuals to receive benefits after exhausting their regular state program benefits. The CARES Act provided an extension of 13 weeks known as Pandemic Emergency Unemployment Compensation (PEUC), scheduled to expire in December 2020. In the Great Recession, federal emergency tiers peaked at 53 weeks, allowing individuals to claim benefits for a total of up to 99 weeks, generating fierce political debate (Chodorow-Reich et al., 2019).

As policy-makers debate whether or how much to extend federal benefits, our results provide some guidance on the number of individuals affected by different extension lengths. For example, in our baseline scenario a sizable number of potentially eligible (i.e. laid off) individuals — more than 4.2 million at the peak — will have unemployment durations beyond the 26 weeks provided by

[7] Fujita and Moscarini (2017) present evidence that *ex post* recall typically exceeds *ex ante* expected recall. We suspect that if anything the opposite will be true in the COVID recession.

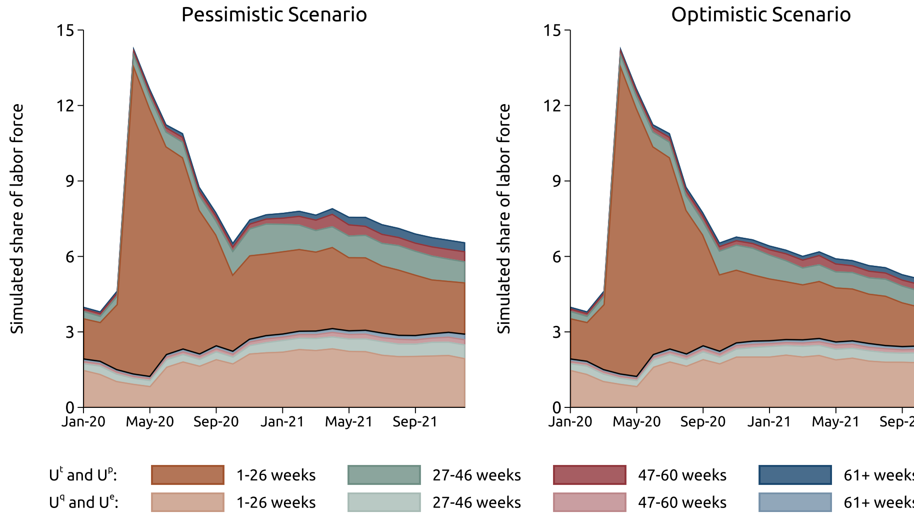

Fig. 5. Alternative Scenarios: Projected Duration Distributions. *Notes:* The figure plots the simulated share of the labor force laid off (U^t and U^p) and other unemployed (U^q and U^e), by the number of weeks in the past two years the individual has been unemployed. The overall path of unemployment follows the November Blue Chip average of the 10 highest unemployment rate forecasts (left panel) or ten lowest unemployment rate forecasts (right panel).

regular state benefits. However, most of these individuals return to employment before 46 weeks, suggesting that they might be covered under the existing EB architecture if that structure were fully utilized.[8] A larger number of individuals remain unemployed for very long durations in the pessimistic scenario, providing a possible rationale for economic-based triggers to govern additional extensions.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at https://doi.org/10.1016/j.jpubeco.2021.104398.

References

Ahn, Hie Joo, Hamilton, James, 2020a. Heterogeneity and unemployment dynamics. J. Bus. Econ. Stat. 38 (3), 554–569.
Ahn, Hie Joo, Hamilton, James, 2020b. Measuring Labor-Force Participation and the Incidence and Duration of Unemployment. Working Paper 27394. National Bureau of Economic Research.
Baker, Michael, 1992. Unemployment duration: compositional effects and cyclical variability. Am. Econ. Rev. 82 (1), 313–321.
Barrero, Jose Maria, Bloom, Nick, Davis, Steven, 2020. COVID-19 Is Also a Reallocation Shock.
Bartik, Alexander, Bertrand, Marianne, Lin, Feng, Rothstein, Jesse, Unrath, Matthew, 2020. Measuring the labor market at the onset of the COVID-19 crisis.
Bewley, Truman, 1980. The optimum quantity of money. In: Kareken, John, Wallace, Neil (eds.), Models of Monetary Economics.
Cajner, Tomaz, Crane, Leland, Decker, Ryan, Grigsby, John, Hamins-Puertolas, Adrian, Hurst, Erik, Kurz, Christopher, Yildirmaz, Ahu, 2020. The U.S. labor market during the beginning of the pandemic recession.
Chodorow-Reich, Gabriel, Coglianese, John, 2019. Unemployment Insurance and Macroe-conomic Stabilization. In: Boushey, Heather, Nunn, Ryan, Shambaugh, Jay, Recession Ready: Fiscal Policies to Stabilize the American Economy, pp. 153–179.
Chodorow-Reich, Gabriel, Coglianese, John, Karabarbounis, Loukas, 2019. The macro effects of unemployment benefit extensions: a measurement error approach. Q. J. Econ. 134 (1), 227–279.
Coglianese, John, 2018. The Rise of In-and-Outs: Declining Labor Force Participation of Prime Age Men.
Drew, Julia A., Rivera, Sarah Flood, Warren, John Robert, 2014. Making full use of the longitudinal design of the current population survey: methods for linking records across 16 months. J. Econ. Soc. Meas. 39 (3), 121–144.
Flood, Sarah, King, Miriam, Rodgers, Renae, Ruggles, Steven, Warren, John Robert, 2020. Integrated Public Use Microdata Series Current Population Survey: Version 7.0.
Fujita, Shigeru, Moscarini, Giuseppe, 2017. Recall and unemployment. Am. Econ. Rev. 107 (12), 3875–3916.
Ganong, Peter, Noel, Pascal, 2019. Consumer spending during unemployment: positive and normative implications. Am. Econ. Rev. 109 (7), 2383–2424.
Hall, Robert E., Kudlyak, Marianna, 2019. Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics. Working Paper 25625. National Bureau of Economic Research.
Jarosch, Gregor, 2015. Searching for Job Security and the Consequences of Job Loss.
Jarosch, Gregor, Pilossoph, Laura, 2018. Statistical discrimination and duration dependence in the job finding rate. Rev. Econ. Stud. 86 (4), 1631–1665.
Kroft, Kory, Lange, Fabian, Notowidigdo, Matthew J., 2013. Duration dependence and labor market conditions: evidence from a field experiment*. Q. J. Econ. 128 (3), 1123–1167.
Kroft, Kory, Lange, Fabian, Notowidigdo, Matthew J., Katz, Lawrence F., 2016. Long-term unemployment and the great recession: the role of composition, duration dependence, and nonparticipation. J. Lab. Econ. 34 (SI), S7–S54.
Krueger, Alan, Cramer, Judd, Cho, David, 2014. Are the Long-Term Unemployed on the Margins of the Labor Market? Brookings Papers on Economic Activity: 229–280.
Krueger, Alan, Mas, Alexandre, Niu, Xiaotong, 2017. The evolution of rotation group bias: will the real unemployment rate please stand up?. Rev. Econ. Stat. 99 (2), 258–264.
Kudlyak, Marianna, Lange, Fabian, 2018. Measuring Heterogeneity in Job Finding Rates among the Non-Employed Using Labor Force Status Histories. Tech. rep.
Mueller, Andreas I., Spinnewijn, Johannes, Topa, Giorgio, 2018. Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence and Bias. Working Paper 25294. National Bureau of Economic Research.
Yagan, Danny, 2019. Employment hysteresis from the great recession. J. Polit. Econ. 127 (5), 2505–2558.

[8] Under existing law states may opt out of EB triggers, and the majority do so, presumably because states must pay for 50% of EB.