

Measuring Labor-Force Participation and the Incidence and Duration of Unemployment

Hie Joo Ahn [‡]
Federal Reserve Board

James D. Hamilton [‡]
University of California, San Diego

March 16, 2019
Revised: June 6, 2020

Abstract

The underlying data from which the U.S. unemployment rate, labor-force participation rate, and duration of unemployment are calculated contain numerous internal contradictions. This paper catalogs these inconsistencies and proposes a unified reconciliation. We find that the usual statistics understate the unemployment rate and the labor-force participation rate by about two percentage points on average and that the bias in the latter has increased over time. The BLS estimate of the average duration of unemployment substantially overstates the true duration of uninterrupted spells of unemployment and misrepresents what happened to average durations during the Great Recession and its recovery.

Keywords: unemployment rate, labor-force participation rate, unemployment duration, measurement error

*The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. We thank Katharine Abraham, Alessandro Barbarino, Travis Berge, Michael Elsby, Andrew Figura, Glenn Follette, Ryan Michaels, Norm Morin, John Stevens and Robert Valletta for comments on earlier drafts of this paper and Jesse Wedewer for excellent research assistance. Data and software to reproduce results in this paper available at http://econweb.ucsd.edu/~jhamilton/AH2_code.zip.

[†]E-mail: hiejoo.ahn@frb.gov

[‡]E-mail: jhamilton@ucsd.edu

1 Introduction.

The Current Population Survey (CPS) is the primary source of information about the labor-force participation rate, unemployment rate, and duration of unemployment for the United States. There are multiple internal inconsistencies in the data from which the fundamental statistics are calculated— if one reported number is correct, another must be wrong. In this paper we catalog these inconsistencies and propose a unified reconciliation of all the problems.

Rotation bias. One source of inconsistency is rotation bias. In any given month, some households are being visited for the first time (rotation 1), others are being interviewed for the second time (rotation 2), with 8 different rotations contributing to the statistics reported for that month. One would think that in a random sample, the numbers calculated from different rotations for a given month should all be the same. But as documented by Hansen et al. (1955), Bailer (1975), Solon (1986), Halpern-Manners and Warren (2012) and Krueger, Mas, and Niu (2017), the reported unemployment rate differs significantly across rotations. In our sample (July 2001 to February 2020), the average unemployment rate among those being interviewed for the first time is 6.6%, whereas the average unemployment rate for the eighth rotation is 5.7%. Even more dramatic is the rotation bias in the labor-force participation rate. This averages 65.8% for rotation 1 and 64.2% for rotation 8 in our sample. Rotation bias affects any inference one draws from the CPS data. For example, it means that if one follows a fixed group of individuals over time, on average outflows from unemployment seem to exceed inflows.

Missing observations. A second source of inconsistency documented by Abowd and Zellner (1985) is that missing observations are not random. Meyer, Mok and Sullivan (2015) noted that households in the CPS have become increasingly less likely to answer surveys or to provide all answers. The standard approach is to calculate statistics for a given month based only on individuals for whom there is an observation that month. But if missing observations are not randomly drawn from the overall population, this may be an increasing source of bias in CPS estimates.

Reported job-search durations and observed continuation probabilities. A third problem in the CPS is inconsistency between the duration of job search reported by an individual in month t and the labor-force status recorded for that same individual in $t - 1$. For example, consider those individuals who were counted as not in the labor force when in rotation 1 in month $t - 1$ and

unemployed when surveyed in rotation 2 in month t . In the second survey, the individual would be asked how long he or she has been looking for work. Two-thirds of these individuals' duration of unemployment is recorded as longer than 4 weeks and 17% of their durations are recorded as one or two years.

A related anomaly is the inconsistency between unemployment hazard rates and the reported duration of unemployment. For example, according to BLS adjusted numbers on labor-force flows, the average unemployed individual in 2011 had a 38% probability of exiting unemployment the following month. Among those already unemployed for more than 6 months, the probability was 31%.¹ From those probabilities we might expect an average duration of unemployment around $(1/0.31)$ or 3 months. Yet according to the BLS, the average duration of unemployment among all those unemployed in 2011 was 40 weeks – three times the value that would be predicted on the basis of the reported hazards.

Number preference. A final source of inconsistency arises from people's preference for reporting certain numbers over others, as documented for example by Baker (1992), Torelli and Trivellato (1993), and Ryu and Slottje (2000). On average there are more people who say they have been looking for work for 6 months than say they have been looking for 23 weeks, though the fraction of those unemployed for 23 weeks should be greater than that of those unemployed for 6 months. In addition, people are more likely to report an even number of weeks than an odd number for shorter spells.

Our proposed resolution. Each of the problems above has been discussed in the literature. Previous papers addressed one problem in isolation. In this paper we show how these problems interact to influence the statistics that economists rely on and propose a unified resolution that addresses the issues simultaneously.

Our first step is to add a fourth category of labor-force status. We regard an individual in any month as either employed (E), unemployed (U), not in the labor force (N), or missing from the sample that month (M). On this basis we construct a data set in which all identities relating stocks and flows are respected; for example, the sum of EE , NE , ME , and UE transitions between $t - 1$ and t exactly equals the total number of E at t . This has never been done before.

¹Our own direct estimates in Panel A of Figure 4 below suggest that hazards do not change much after durations beyond six months.

Our second step is to model statistically the way in which people’s answers change the more times they have been interviewed. We interpret households in different rotations as being surveyed using a different interview technology and summarize how the differences in the average answers given by different rotations change gradually over time. We calculate the answer to the following counterfactual question: if a group of households in rotation j in month t were being interviewed for the first time instead of the j th time, how would their answers have been different?

Combining these two steps allows us to produce the first fully reconciled description of stocks and flows in the CPS data. By then looking at how reconciled ME , MN , and MU transitions differ from the rest of the population, we are able to adjust the treatment of missing observations based on what we know about those individuals when data are collected from them.² We find that missing individuals are more likely than the general population to be unemployed. In addition, the biases introduced by missing observations have increased over time and are bigger when the labor market is slack. Our paper is the first to document the cyclical features in the bias coming from nonrandom missing observations.

With these tools we can provide the first-ever reconciled description of stocks and flows normalized on the basis of any of the eight interview technologies. In practice we need to choose a particular technology, which requires taking a stand on the source of the rotation bias. Much evidence, most of which is new to this paper, persuades us that the primary source of rotation bias is that individuals who are unemployed or only employed part-time perceive some stigma associated with those answers and become less engaged with the interview process the more times an interview is attempted. We conclude that the first-interview concept of unemployment is the best one to use, and show that the unemployment rate and labor-force participation rate are more seriously underestimated at points in time or for demographic groups for which the reported unemployment rate is higher.

Reconciling the inconsistencies in the CPS further requires confronting the gray area in the distinction between unemployed and not in the labor force. We argue that the most satisfactory way

²Our adjustment for missing observations is similar in spirit to that in Abowd and Zellner (1985), though there are a number of important differences. For any month t they make one adjustment looking backward in time and a second adjustment looking forward in time, giving them potentially two different estimates for each month t . By contrast, we take a unified approach to the full data set. Abowd and Zellner’s adjustments do not deal with the problem of rotation bias or the other measurement issues for which we also develop solutions. And they only calculate average unemployment rates over what is now a historically old sample. By contrast, we adjust estimates month-by-month up to the present.

to resolve the inconsistencies is to reclassify some of those currently counted as N as instead actively seeking work. As noted above, in every month t a significant fraction of Americans are counted as N_t but the following month are counted as U_{t+1} and report at $t+1$ that they have been looking for work for 5 weeks or longer. We find that the incidence of such reports is explained not by education or demographic group but instead can be very accurately predicted by the reported incidence of medium- and long-term unemployment at date t . The incidence can separately be predicted by whether those individuals said at time t that they wanted a job. Moreover, individuals in this category have similar future job-finding probabilities and unemployment-continuation probabilities as if they had been counted as U_t . Our proposal is therefore to classify such individuals as U_t . This is another reason why the unemployment rate and labor-force participation rate are more seriously underestimated at points in time or for demographic groups for which the unemployment rate is higher.

Finally, our resolution to the number-preference problem is to postulate a flexible latent distribution of perceived durations that is then reported by individuals with a certain structure of number-reporting preference; for related approaches see Baker (1992), Torelli and Trivellato (1993), and Ryu and Slottje (2000). Our approach is completely new compared to these studies in that our parameterization allows direct linkage of data on stocks, flows, and durations and in that both digit and interval preference are jointly considered. Our framework describes the reported values extremely accurately.

Why does it matter? The importance of these issues is illustrated in Panel A of Figure 1. This asks a very fundamental question: if someone is unemployed at $t-1$, what is the probability that person will still be unemployed at t ? Researchers have used the CPS data to answer this question in two different ways. A measure based on reported unemployment durations calculates the ratio of individuals who are unemployed at t with a reported duration greater than 4 weeks to the total number of individuals unemployed at $t-1$. Variants of this calculation have been used by van den Berg and van der Klaauw (2001), Elsby, Michaels and Solon (2009) and Shimer (2012). This measure is plotted as the black line in Panel A. An alternative measure based on labor-force flows looks at the subset of individuals who are U at $t-1$ and either E , N , or U at t and calculates the number of UU continuations as a fraction of the sum. Variants of this approach were used by Fujita and Ramey (2009) and Elsby, Hobijn and Şahin (2010). The flow-based measure is plotted

as the green line. If all magnitudes were measured accurately the two estimates should give a similar answer. But in practice they are wildly different. The duration-based measure averages 68.6% over our sample, while the flow-based measure averages 53.0%.

These differences are caused by the multiple inconsistencies mentioned above. The flow-based measure underestimates the true continuation probability because (1) some *UN* transitions are a result of rotation bias and (2) some *UN* transitions should be interpreted as *UU* continuations. The duration-based measure overestimates the probability, because a substantial number of people interpret the duration of job search as including on-the-job search or the time since the last salient job; see Elsby et al. (2011), Farber and Valletta (2015), and Kudlyak and Lange (2018). Our reconciled estimate is shown in the blue line in Panel A, and falls in between the other two estimates. The flow-based estimate was closer to our adjusted measure at the beginning of the sample, whereas the duration-based measure is closer to our adjusted series today.

Another fundamental question is, how many people become unemployed each month? One estimate (e.g., Shimer, 2012) is simply the number of unemployed individuals reporting durations of less than 5 weeks. The black line in Panel B of Figure 1 shows this value as a percent of the civilian noninstitutional population. As noted by Elsby et al. (2011), it underestimates new inflows into unemployment since half of *EU* and *NU* transitions report unemployment durations of 5 weeks or longer. Alternatively, the BLS publishes separate estimates of *EU* and *NU* flows that they adjust to address some of the problems that we document in this paper. However, our analysis suggests that rotation bias and non-random missing observations have not yet been completely corrected for in the BLS adjusted series (in green). Our reconciled series (blue) is often significantly higher than the BLS adjusted estimate.

Panel C of Figure 1 compares our adjusted estimate of the unemployment rate with the BLS estimate. Our measure is 2.0% higher on average, and the gap increased during the Great Recession. The gap recovered gradually after the recession and has only recently returned to its pre-recession level. The gap between our measure and the BLS measure of the labor-force participation rate (Panel D) is 2.1% on average. It also increased in the Great Recession and remained elevated as of 2020. We conclude that labor-force participation declined slightly less over this period than suggested by the BLS series.

Whereas BLS estimates of unemployment duration are based on individuals' reported durations

of job search, our estimates are based on uninterrupted spells of unemployment. In going from the green to blue lines in Panel A, we adjusted unemployment continuations up considerably from the standard estimates, but we did not adjust these all the way up to those implied by reported durations in black. As a result, our reconciled estimates of average unemployment durations (shown as blue in Panel E) are considerably below those from BLS (black), similar to the conclusion by Kudlyak and Lange (2018). Our estimates of average duration did not rise as much during the Great Recession as suggested by the BLS series based on reported durations. Also, our reconciled estimates subsequently recovered to pre-recession levels, whereas the BLS reported durations do not.

A significant part of the measurement errors we discuss arises from ambiguities in classifying individuals as “unemployed” versus “not in the labor force.” The employment-to-population ratio (panel F) avoids these issues and thus might be a better indicator of labor market slack. However, the employment-to-population ratio is still influenced by rotation bias, which we attribute to stigma and disengagement with repeated interviews of some part-time workers.

A number of important studies have approached the problem of measurement error in the CPS data in a very different way from ours. A common assumption is that the reported data differ from latent true values, with identification coming from assumptions about the joint dynamics of the true values and measurement error. Prominent examples include Biemer and Bushery (2000), Feng and Hu (2013), and Shibata (2019). These studies are silent about the source of misclassification errors, and did not deal with rotation bias, nonrandom missing observations, inconsistency between reported duration and the previous labor force status, or reporting errors of unemployment duration in their correction. Biemer and Bushery (2000) and Shibata (2019) assumed that true labor-force transitions were first-order Markov. Feng and Hu (2013) relaxed this assumption, though Shibata (2019) concluded that their approach generates implausible transition probabilities. By contrast, our approach does not impose any Markov assumptions and produces plausible transition probabilities and unemployment durations that are consistent with these probabilities. Our approach also explains well the non-Markov predictability of labor-force status documented by Kudlyak and Lange (2018). Although our methods and assumptions are very different from these studies, we nevertheless reach a similar conclusion that the BLS significantly underestimates the average unemployment rate and overestimates the average duration of unemployment.

The plan of the paper is as follows. Section 2 describes the structure of the CPS survey and our data set. Section 3 uses averages over the complete sample to document the various inconsistencies in the CPS data and develops the statistical framework that will form the basis for our reconciliation. Section 4 describes the steps we use to reconcile these inconsistencies. Section 5 briefly concludes.

2 Data construction.

Since July 2001, each month around 60,000 housing units are included in the Current Population Survey. An effort is made to contact each address and determine the number of individuals aged 16 or over who are not in the armed forces or in an institution such as prison or a nursing home. An individual is counted as employed (E) if during the reference week of the survey month the individual did any work at all for pay, for their own business, or were temporarily absent from work due to factors like vacations, illness, or weather. People are counted as unemployed (U) if they were not E but were available for work and made specific efforts to find employment some time during the previous 4 weeks. Individuals who are neither E nor U are counted as not in the labor force (N). One person in the household can provide separate answers for each of the individuals living at that address.

The next month and each of the following two months, the interviewer attempts to contact the same address to ask the same questions. In any given month, around 1/8 of the 60,000 qualifying households being interviewed for the first time (denoted rotation 1), and another 1/8 each are being interviewed for the second, third or fourth time (rotations 2, 3, or 4). After the fourth month the household is not interviewed for the next 8 months, but is reinterviewed again 1 year after the first interview (rotation 5) and again for each of the following 3 months (rotations 6, 7, and 8). For data since 1994, if an individual was unemployed in two consecutive months, the interviewer does not ask again the duration of unemployment the second month, but simply adds time elapsed since the previous interview to the previous answer. Thus new unemployment duration data are only collected in rotations 1 and 5, or in the other rotations for someone who was E , N or missing from the sample the month before.

The survey is imperfect for purposes of tracking the experience of an individual across months

due to various measurement problems; for discussion of these see Madrian and Lefgren (2000) and Nekarda (2009). Each address has a unique identifier, and an effort is made to associate an individual person within that household with a particular 2-digit number. Our study is unique in treating missing (M) as a separate observed category for someone whose information is not available in a particular rotation or is inconsistent from the information reported for that individual in other rotations. As in Abowd and Zellner (1985), we will use information about that individual in months where it is available to correct for the fact that individuals who are sometimes missing (which could come in part from households that are more prone to reporting errors or to having people moving in or out) may differ in systematic ways from individuals for whom 8 separate months of data are available. We check if an individual with the same household and personal identifier is reported to have the same gender and an age that does not differ by more than 2 years across rotations. If so, we consider that individual successfully matched. If not, we designate that individual as M in the months for which no status is available or for which the age and gender records are inconsistent with those reported across the majority of the 8 rotations.³

The raw data for our study thus consist of $y_{X,t}^{[j]}$, the sum of the number of individuals (multiplied by a weight associated with that individual) who are in rotation $j \in \{1, \dots, 8\}$ in month t with reported status $X \in \{E, N, M, U\}$, and $y_{X_1, X_2, t}^{[j]}$, the weighted sum of individuals reporting X_1 in rotation $j - 1$ in month $t - 1$ and X_2 in rotation j in month t for $j \in J = \{2, 3, 4\} \cup \{6, 7, 8\}$. See Table A-1 in the online appendix for a summary of notation used in this study. A key advantage of our approach is that, unlike the values used by any other researchers, our data on stocks and flows are internally consistent by construction, always satisfying the accounting identities

$$y_{X_2, t}^{[j]} = y_{E, X_2, t}^{[j]} + y_{N, X_2, t}^{[j]} + y_{M, X_2, t}^{[j]} + y_{U, X_2, t}^{[j]} \quad (1)$$

$$y_{X_1, t-1}^{[j-1]} = y_{X_1, E, t}^{[j]} + y_{X_1, N, t}^{[j]} + y_{X_1, M, t}^{[j]} + y_{X_1, U, t}^{[j]} \quad (2)$$

³Nekarda (2009) used race in addition to age and gender and Madrian and Lefgren (2000) also used education. Both these variables are susceptible to ambiguity and could be reported differently for a fixed individual, particularly when a different individual answers the questions for the household. We topcode age at 65 years or older, so an individual in this age group with the same address, same gender, and same identifying number is considered matched. Note that our category of M includes people who move into the address (for example, from another address, getting out of prison or the military, or becoming 16 years old) and people who leave the address (whether to another address, institution or through death). Our M category also includes people who do not answer the questions in some rotations or who answer the gender and age questions in an inconsistent way.

for all t, X_1, X_2 and $j \in J$.

One drawback of this procedure is that we need 16 months of observations to determine whether to categorize someone as M in a given month. For example, our sample starts in 2001:7. Someone whose history beginning in 2001:7 was $EEMM - MMMM$ will be counted as M in rotation 3 in 2001:9 by our method, whereas someone who would have had the same history if initially surveyed in 2001:5 would never appear in the sample.⁴ This causes the number of individuals who are classified as M to be artificially depressed in the first year of the sample. A similar effect arises at the end of the sample, with individuals whose record would have been $MMEE - EEEE$ not being apparent if their rotations 1 or 2 would have come at the end of the sample. We therefore adjusted the counts of M and MM at the beginning and end of the sample upward based on the average counts of M for each rotation over the nearest year of complete observations; for details see Appendix A. Since changes in M occur relatively slowly in our sample, this adjustment has little effect on any of the key measures we develop. We made additional adjustments when new households were added and other households dropped in the 2004 and 2014 sample redesigns.⁵

BLS also assigns a weight to each individual. People with characteristics that are underrepresented in a particular month are given a larger weight. These weights are a partial response of BLS to the issue that missing individuals are not a random sample of the population. We want to include this correction to demonstrate the need for additional corrections for missing individuals. We can not use the exact BLS weights to do this because the BLS may assign a given individual different weights in two different months, which is another reason in addition to missing observations why (1) and (2) do not hold in the BLS data. Our approach was to assign a fixed weight for an individual across all 8 possible observations based on the BLS weight for that individual in the first month for which data are recorded for that person, as described in Appendix A.

3 Statistical description of labor-force status data.

In this section we develop statistical descriptions of a number of features of the CPS data.

⁴See Appendix A for detailed examples.

⁵With the expansion of the survey from 50,000 to 60,000 households, beginning in July 2001, some individuals were added and others dropped across a number of rotations, with waves of new individuals added to subsequent rotations 5. Tracking individuals before and after this break is considerably harder than handling the sample redesign in 2004 and 2014. For this reason we simply begin our analysis with the modern design adopted in July 2001.

3.1 Unemployment durations reported in rotations 1 and 5.

First we consider the durations of unemployment that are reported on average over our sample by people who are being interviewed for the first time (rotation 1). The blue bars in the top panel of Figure 2 plot the fraction of unemployed reporting the indicated duration of job search in weeks. Clearly there are some significant reporting errors arising from number preference. Respondents are more likely to report spells as an integer number of months, and for longer spells as either 6 months, 1 year, 18 months, or longer than 99 weeks. For shorter spells, people are more likely to report an even number of weeks instead of an odd number; for example, on average there are more people reporting 2 weeks than 1 and 6 weeks than 5. Respondents are extremely unlikely to report a duration of zero weeks, and for this reason we group the 0-week and 1-week observations together into a category of reported duration less than or equal to one week.

To interpret these numbers in an internally consistent way, we impose the restriction that the only way an individual could have been unemployed for τ weeks would be if the individual had been unemployed for $\tau - 1$ weeks the week before. Thus if $\pi_U^\dagger(\tau)$ denotes an internally consistent summary of the fraction of the population who have been searching for τ weeks, the function $\pi_U^\dagger(\tau)$ must be monotonically decreasing in τ . For our baseline specification we propose to represent this function as a mixture of two exponentials with decay rates p_1 and p_2 , respectively. We form a (99×1) vector π_U^\dagger whose τ th element for $\tau = 1, 2, \dots, 98$ is an internally consistent representation of the fraction of the working-age population who perceive having been unemployed for a duration of τ weeks at a fixed point in time, while the 99th element is the fraction with perceived duration greater than 98 weeks:

$$\pi_U^\dagger = \pi_{1U}^\dagger + \pi_{2U}^\dagger \quad (3)$$

$$\pi_{iU}^\dagger = \pi_U w_i (1 - p_i) \begin{bmatrix} 1 & p_i & p_i^2 & \cdots & p_i^{97} & p_i^{98}/(1 - p_i) \end{bmatrix}' \quad \text{for } i = 1, 2. \quad (4)$$

Here π_U denotes the fraction of the population who are unemployed and w_i the fraction of those individuals who are type i . Such a distribution would be the outcome of a steady state in which there was a fraction $\pi_U w_1 (1 - p_1)$ of the population who lose their jobs each week and for each of whom the probability of continuing unemployed in any subsequent week is p_1 , and an additional

inflow of $\pi_U w_2(1 - p_2)$ individuals with continuation probability p_2 .⁶

We allow for the various forms of number preference noted above by introducing a (99×99) matrix $A(\theta_A)$ whose elements are determined by a (13×1) vector θ_A . The first element $\theta_{A,1}$ allows a preference for reporting short durations as an even rather than an odd number of weeks, assuming that someone whose true duration is $\tau = 1, 3, 5$, or 7 in fact reports duration $2, 4, 6$, or 8 with probability $\theta_{A,1}$. The value of $\theta_{A,2}$ represents the probability that someone will round their duration up or down by a week to reach an integer number of months for durations within one week of $1, 2, 3$ or 4 months, while someone two weeks away from either of two months is presumed to round down with probability $\theta_{A,3}/2$ and up with probability $\theta_{A,3}/2$. As we move to longer durations we allow for the possibility that the rounding tendencies become stronger, introducing new pairs of parameters for durations between 5 - 7 months, 8 - 11 months, or 12 or more months. The last elements of θ_A allow for preferences for integer multiples of 6 months for longer durations. For each τ the τ th column of A sums to unity and characterizes the probability that someone whose true duration category is τ will report each of the possible categories i between 1 and 99 , where i or $\tau = 99$ is interpreted as true or reported durations longer than 98 weeks. Appendix B provides more details on the structure we use to represent the matrix A . Note that our framework does not impose the assumption of the existence or magnitude of any particular reporting error, as it includes as a special case no reporting error of any kind when $\theta_A = 0$.

Let $y_{X,t}^{[1]}$ be the number of individuals in rotation group 1 sampled at date t who report status X for X one of E (employed), N (not in labor force), M (labor-force status for that individual is missing), or U (unemployed). We summarize further detail in the last category in terms of $y_{U,t}^{[1]}(\tau)$ which is the number of unemployed who report having been looking for work for τ weeks for $\tau = 1, \dots, 99$.⁷ We compare the observed values $y_{U,t}^{[1]}(\tau)$ with the predicted values represented by the (99×1) vector

$$\dot{\pi}_U = A\pi_U^\dagger. \quad (5)$$

We also let π_X denote the overall fraction of the population reporting status $X \in \{E, N, M, U\}$.

⁶We will later examine some testable implications of such an interpretation by looking at the actual unemployment-continuation probabilities for different individuals and also look at alternative functional forms. But for now we propose (3) and (4) as a simple but flexible parametric functional form with which to impose monotonicity on $\pi_U^\dagger(\tau)$.

⁷The duration is top-coded at 99 weeks in our data.

If we treated observations as independent across months t the log likelihood of the rotation 1 observations alone would then be

$$\begin{aligned}\ell_X^{[1]}(\lambda_X) &= \sum_{t=1}^T [y_{E,t}^{[1]} \ln \pi_E + y_{N,t}^{[1]} \ln \pi_N + y_{M,t}^{[1]} \ln \pi_M] \\ &\quad + \sum_{t=1}^T \sum_{\tau=1}^{99} y_{U,t}^{[1]}(\tau) \ln \dot{\pi}_U(\tau).\end{aligned}\tag{6}$$

We can maximize this with respect to $\theta_A, p_1, p_2, w_1, w_2, \pi_E, \pi_N, \pi_M, \pi_U$ subject to the constraint that all probabilities are between 0 and 1 and sum to unity.⁸

Estimates are reported in column 1 of Table 1, along with quasi-maximum-likelihood standard errors in column 2 which allow for the possibility that $y_{X,t}^{[1]}$ is correlated across time (calculated as described in Appendix C). The predicted reported values $\dot{\pi}_U(\tau)$ are compared with the average reported values in the top panel of Figure 2.⁹ This framework is able to describe the reported values extremely accurately. The estimated latent function $\pi_U^\dagger(\tau)$ along with its two contributing components are plotted as a function of τ in the bottom panel of Figure 2. We also considered an alternative functional form based on a Weibull distribution, as discussed in Appendix D. The mixture of exponentials has a much better fit to the data than that for the Weibull specification, and we will use it in our baseline analysis.

For rotations 2-4 and 6-8, BLS imputes a duration to those reporting UU continuations, making durations for these individuals a hybrid of perceived and imputed quantities. This can create a downward bias in the number of individuals unemployed for less than 5 weeks as discussed by Abraham and Shimer (2001) and Shimer (2012) and blurs the inconsistency between perceived and imputed durations. Since our goal is to characterize perceived durations separately from objective durations, we do not use the imputed duration in the second month in unemployment. However, there are no imputations for unemployment durations for those people in rotation 5. We therefore repeated the analysis with $y_{X,t}^{[1]}$ in (6) replaced by $y_{X,t}^{[5]}$. Parameter estimates and standard errors are reported in columns 3 and 4 of Table 1. These are very similar to those inferred from the

⁸Maximum likelihood estimates of some parameters are known analytically. Let $y_X = \sum_{t=1}^T y_{X,t}$ denote the total number of observations in category X and $n = (y_E + y_N + y_M + y_U)$ the total number of observations. Then $\hat{\pi}_X = y_X/n$ for $X \in \{E, N, M, U\}$. These values can be substituted into expression (6) and the resulting concentrated likelihood then maximized with respect to θ_A, p_1, p_2, w_1 with $w_2 = 1 - w_1$.

⁹As noted in the previous footnote, by the nature of the maximization problem, the estimated values $\hat{\pi}_X$ for $X = E, N, M$ exactly match the historical fractions $y_X/(y_E + y_N + y_M + y_U)$.

rotation 1 observations alone.

3.2 Characteristics of NU , EU , and MU transitions.

Next consider the status of individuals in rotation 2 who had been counted as not in the labor force when surveyed in rotation 1. Figure 3 focuses on the subset who in the second month (when they were in rotation 2) reported being unemployed, giving the percentage reporting each duration of job search. Two-thirds of these people have a duration of unemployment in rotation 2 that is recorded to be longer than 4 weeks, despite the fact that the previous month they did not report actively looking for a job and so were counted as out of the labor force. 8.6% of NU individuals say that they have been looking for a job for 52 weeks and another 8.5% report having been looking for work for over 98 weeks.

Of those people who report right after an NU transition that they have been looking for work for more than 4 weeks, what distribution characterizes their perceived duration of job search? We represent the probability of transitions from N to E, N, M , or U with parameters $\pi_{NE}, \pi_{NN}, \pi_{NM}, \pi_{NU}$, respectively, where these four numbers sum to unity. Of those who make an NU transition and report an unemployment duration greater than 4 weeks, suppose that their perceived duration can again be represented by a mixture of two exponentials with decay parameters $p_{1,NU}$ or $p_{2,NU}$. We assume that some fractions $q_{1,NU}, q_{2,NU}, q_{3,NU}$, and $q_{4,NU}$ of those making the NU transition will perceive their unemployment duration to be 1, 2, 3, or 4 weeks respectively, treating these values of $q_{j,NU}$ completely unrestrained. A fraction $q_{5,NU}$ perceive a duration greater than 4 weeks drawn from an exponential distribution with parameter $p_{1,NU}$ and a fraction $q_{6,NU}$ are characterized by $p_{2,NU}$, with $\sum_{j=1}^6 q_{j,NU} = 1$. We thus calculate

$$\pi_{NU}^{\dagger}(\tau) = \begin{cases} q_{\tau,NU} & \text{for } \tau = 1, 2, 3, 4 \\ q_{5,NU}(1 - p_{1,NU})p_{1,NU}^{\tau-5} + q_{6,NU}(1 - p_{2,NU})p_{2,NU}^{\tau-5} & \text{for } \tau = 5, 6, \dots, 98 \\ q_{5,NU}p_{1,NU}^{94} + q_{6,NU}p_{2,NU}^{94} & \text{for } \tau = 99 \end{cases} \quad (7)$$

The predicted probability of each reported duration is then given by $\dot{\pi}_{NU} = \pi_{NU} A \pi_{NU}^{\dagger}$.

Let $y_{NX,t}^{[2]}$ denote the number of individuals who counted as not in the labor force in rotation 1 in month $t - 1$ and reported status X at date t where $X \in \{E, N, M, U\}$. Let $y_{NU,t}^{[2]}(\tau)$ denote the

number of NU who report unemployment duration $\tau \in \{1, \dots, 98, \geq 99\}$ in rotation 2. Then the contribution to the likelihood for months $t = 1, \dots, T$ from rotation 2 NX transitions is

$$\begin{aligned} \ell_{NX}^{[2]}(\lambda_{NX}) = & \sum_{t=1}^T [y_{NE,t}^{[2]} \ln \pi_{NE} + y_{NN,t}^{[2]} \ln \pi_{NN} + y_{NM,t}^{[2]} \ln \pi_{NM}] \\ & + \sum_{t=1}^T \sum_{\tau=1}^{99} y_{NU,t}^{[2]}(\tau) \ln \hat{\pi}_{NU}(\tau). \end{aligned} \quad (8)$$

This expression can then be maximized with respect to $\lambda_{NX} = (\theta'_{A,NU}, p_{1,NU}, p_{2,NU}, \pi_{NE}, \pi_{NN}, \pi_{NM}, \pi_{NU}, q_{1,NU}, q_{2,NU}, \dots, q_{6,NU})'$ subject to the constraints that all parameters fall between 0 and 1, $\pi_{NE} + \pi_{NN} + \pi_{NM} + \pi_{NU} = 1$ and $\sum_{j=1}^6 q_{j,NU} = 1$.

Quasi-maximum-likelihood estimates $\hat{\lambda}_{NX}$ are reported in column 5 of Table 1 and predicted values $\hat{\pi}_{NU}$ compared with historical average values for y_{NU} in Figure 3. Note that θ_A was estimated in column 1 solely from individuals who were recorded as being unemployed in rotation 1, in column 3 solely from individuals who were unemployed in rotation 5, and in column 5 solely from individuals who were recorded as being out of the labor force in rotation 1 and unemployed in rotation 2. Although the vector θ_A was estimated from very different data, the estimated values are quite similar. Likewise $\hat{p}_{1,NU}$ and $\hat{p}_{2,NU}$ turn out to be very close to the values \hat{p}_1 and \hat{p}_2 estimated from rotations 1 and 5. Those who make NU^{5+} transitions, like those who are unemployed in rotation 1, are allowed to answer any number to the question, “how long had you been looking for a job.” Due to this feature, the question reveals the perceived job-search spells of an individual in a way that the assigned durations for UU individuals does not. The similarity in the parameter estimates suggests that the perceived job-search history of NU^{5+} individuals is similar to that of the pool of unemployed in rotation 1.

Next consider the status in month t of individuals who were recorded as employed when sampled in rotation 1 in month $t - 1$. Twenty-six percent of those who make EU transitions report durations longer than 4 weeks. Unlike the NU transitions, we do not interpret these as necessarily implying an inaccuracy in either the E or U designation. Kudlyak and Lange (2018) noted these could represent records of individuals who were employed in $t - 1$ but were engaged in on-the-job search for a new job.¹⁰ We repeated the procedure to characterize transitions from employment using the

¹⁰Elsby et al. (2011) and Farber and Valletta (2015) suggested that EU^{5+} individuals could also be reporting the time since the last salient job. Both this interpretation, as well as that of Kudlyak and Lange (2018), support the conclusion that the reported duration associated with an EU transition should not be interpreted as the duration of an

same framework as above, replacing NX in (8) with EX . Parameter estimates and standard errors are reported in columns 7 and 8 of Table 1. Much fewer EU transitions perceive themselves as long-time job seekers ($q_{6,EU} = 0.17$ versus $q_{6,NU} = 0.50$). We also looked at the status in rotation 2 of individuals who were missing in rotation 1, replacing EX with MX . Quasi-maximum-likelihood estimates are reported in column 9 of Table 1. Individuals making MU transitions look similar to the pool of unemployed in rotation 1.

3.3 Characteristics of UX transitions.

We next examine UX transitions. The bars in the top panel of Figure 4 show $\hat{\pi}_U(\tau)$, the observed probability that someone in rotation 1 who reports being unemployed with duration τ weeks will still be unemployed the following month.¹¹ This probability rises as a function of duration before eventually plateauing at a value around 0.62 for durations over half a year. One way this feature of the data is often captured is by defining some arbitrary cutoff K with any duration $\tau \leq K$ designated as short-term unemployed who have some continuation probability $\gamma_{1,UU}$ while long-term unemployed ($\tau > K$) have a different probability $\gamma_{2,UU}$. That kind of simple dichotomization into short-term and long-term unemployment would have the drawbacks that it requires picking an arbitrary cut-off K and implies an abrupt discontinuity in outcomes expected for individuals slightly below K relative to those slightly above K .

Our parameterization suggests a smooth function that could be used as a natural alternative to an arbitrary cutoff. We have summarized the distribution of reported durations for those unemployed in rotation 1 as coming from a mixture of two types of individuals, where type 1 have a perceived weekly continuation probability of p_1 and type 2 have a perceived continuation probability of p_2 . We modeled the fraction of the population that reports being unemployed with duration τ as given by the τ th element of the vector $\xi_1 + \xi_2$ where $\xi_i = A\pi_{iU}^\dagger$ for π_{iU}^\dagger given in (4). If we observe someone reports a duration of τ , the probability that the individual is type i is obtained

uninterrupted spell of unemployment. Ahn and Shao (2017) further documented that on-the-job search constitutes a non-negligible fraction of aggregate job search. Hall and Kudlyak (2019) found that many job losers make frequent transitions between short-term employment, unemployment, and out of the labor force before finding a long-term job.

¹¹To avoid plotting values for observations with excessive sampling error, we set this probability to 0 for durations with 10 or fewer observations over the whole sample.

from the formula

$$\eta_i(\tau) = \xi_i(\tau) / [\xi_1(\tau) + \xi_2(\tau)] \quad (9)$$

for $i = 1$ or 2 . The function $\eta_2(\tau)$ is plotted in the bottom panel of Figure 4.¹² Someone who reports a duration of $\tau = 1$ week is quite unlikely to have come from the second distribution, whereas someone who reports a duration greater than 40 weeks is almost certain to have come from the second distribution. The function dips down at duration $\tau = 26$ weeks because, given the tendency of answers to clump at this value, this observation includes many individuals whose true duration is less than 26 weeks and accordingly contains a higher mix of type 1 relative to those reporting 25 weeks.

This formula allows us to estimate objective monthly transition probabilities for the two types. Let $\gamma_{i,UX}$ be the probability that an individual of type i makes a transition from unemployment in rotation 1 to status $X = E, N, M$, or U in rotation 2, so $\gamma_{i,UE} + \gamma_{i,UN} + \gamma_{i,UM} + \gamma_{i,UU} = 1$ for both $i = 1$ and $i = 2$. Let η_i denote the vector whose τ th element is $\eta_i(\tau)$ and $\dot{\pi}_{UX}$ the (99×1) vector whose τ th element is the observed probability that someone who reports duration τ in month t has status X in month $t + 1$. Under the above assumptions $\dot{\pi}_{UX}$ would be predicted to be

$$\dot{\pi}_{UX} = \eta_1 \gamma_{1,UX} + \eta_2 \gamma_{2,UX}. \quad (10)$$

Let $y_{UX,t}^{[2]}(\tau)$ denote the observed number of individuals who report U with duration τ in rotation 1 and status X in rotation 2. We then have the likelihood function

$$\begin{aligned} \ell_{UX}^{[2]}(\lambda_{UX}) = \sum_{t=1}^T \sum_{\tau=1}^{99} & [y_{UE,t}^{[2]}(\tau) \ln \dot{\pi}_{UE}(\tau) + y_{UN,t}^{[2]}(\tau) \ln \dot{\pi}_{UN}(\tau) \\ & + y_{UM,t}^{[2]}(\tau) \ln \dot{\pi}_{UM}(\tau) + y_{UU,t}^{[2]}(\tau) \ln \dot{\pi}_{UU}(\tau)]. \end{aligned} \quad (11)$$

We fixed η_2 to be the function plotted in the bottom panel of Figure 4 and maximized (11) with respect to $\{\gamma_{i,UE}, \gamma_{i,UN}, \gamma_{i,UM}, \gamma_{i,UU}\}_{i=1,2}$ subject to the constraint that $\gamma_{i,UE} + \gamma_{i,UN} + \gamma_{i,UM} + \gamma_{i,UU} = 1$ for $i = 1, 2$.

Quasi-maximum-likelihood estimates and standard errors are reported in rows 1 and 2 of Table 2.

¹²For purposes of this graph, this function was calculated using the values of w_1, p_1, p_2, θ_A from Table 3, which pool all observations from all rotations to estimate these parameters.

Type 1 individuals have a 32% probability of being employed next month, whereas the probability for type 2 individuals is only 12%. Type 1 individuals have a 37% probability of being unemployed next month, whereas for type 2 the probability is 57%. The red line in the top panel of Figure 4 show the predicted values for the unemployment-continuation probability implied by these maximum likelihood estimates.¹³ This function provides a very good summary of the raw data.

We also repeated the analysis using only data for individuals who were unemployed in rotation 5, with very similar results shown in rows 3 and 4. Our preferred estimates pool together all observations for all rotations but still estimate $\gamma_{i,UU}$ completely independently of the value of p_i , while treating the values of θ_A , p_1 , and p_2 as the same across all rotation groups. This summary of the full data set was obtained by maximizing the full-sample likelihood

$$\ell = \ell_X^{[1]} + \ell_X^{[5]} + \sum_{j \in J} \left(\ell_{EX}^{[j]} + \ell_{NX}^{[j]} + \ell_{MX}^{[j]} \right) + \ell_{UX}^{[2]} + \ell_{UX}^{[6]}. \quad (12)$$

These full-sample estimates are reported in Table 3.

Now let us compare the estimated objective unemployment-continuation probability for type 1 individuals ($\gamma_{1,UU}$) with the value that would be predicted on the basis of their reported durations. If type 1 individuals truly had a weekly unemployment-continuation probability of $p_1 = 0.8092$, we would expect to observe a monthly continuation probability of $0.8092^{4.33} = 0.40$. If we condition on missing observations having the same distribution as observed E , N and U , this value turns out to exactly equal the value we'd predict from Table 3 of $\gamma_{1,UU}/(1 - \gamma_{1,UM}) = 0.40$. Note that our approach did not impose this in any way; \hat{p}_1 is based solely on reported durations, whereas $\hat{\gamma}_{1,UU}$ is based solely on observed continuations. The exercise shows that the durations reported by type 1 individuals are entirely consistent with the observed labor-force flows for those individuals.

By contrast, the long-term unemployed are another story. Their perceived weekly unemployment-continuation probability of $p_2 = 0.9733$ would imply a monthly continuation probability of $0.9733^{4.33} = 0.89$, far larger than the estimate $\gamma_{2,UU}/(1 - \gamma_{2,UM}) = 0.62$. Even more dramatically, a monthly continuation probability of 0.62 would mean a probability of remaining unemployed for 6 months of $0.62^6 = 0.06$. But in the BLS data, the fraction of those unemployed who report durations over 26 weeks averages 28%. Far fewer people than are reported in the data

¹³That is, the red line plots $\eta_1(\tau)\hat{\gamma}_{1,UU}/(1 - \hat{\gamma}_{1,UM}) + \eta_2(\tau)\hat{\gamma}_{2,UU}/(1 - \hat{\gamma}_{2,UM})$ as a function of τ .

should be unemployed longer than 6 months if people left the pool of long-term unemployed at anything like the rate implied by $\gamma_{2,UU}$. The observed unemployment continuation probabilities are not consistent with the distribution of reported unemployment durations.

That conclusion is robust whether one uses our parametric model or any other. For example, Appendix D derives the analogous result using a Weibull characterization of durations. Any model that accurately describes the cross-section of durations – and ours does so quite well – will predict an unemployment-continuation similar to the stock-based measure plotted as the black line in Figure 1, which we noted is inconsistent with flow-based measures in green. The main advantage of our parametric approach is that it highlights that this inconsistency between the stock-based and flow-based measures comes entirely from those whom we have characterized as the perceived long-term unemployed. This insight is new to this literature.

3.4 Rotation-group bias.

Another source of error in the CPS data is the difference across different rotations in the reported labor-force status. Table 4 reports the monthly average number of sampled individuals with measured labor force status E, N, M , or U for each of the 8 rotation groups.¹⁴ Column 6 shows that the average unemployment rate declines sharply as a function of rotation group, starting out at 6.6% for rotation 1 but falling all the way to 5.7% for rotation 8. Column 7 reveals another interesting fact that has not been much commented on in the earlier literature: the measured labor-force participation rate falls even more sharply. Column 3 documents a third tendency—individuals are much more likely to be missed in rotation 1 and 5 compared to other groups.

We summarize these tendencies with some simple regressions. Let $x_t^{[j]} = 100y_{X,t}^{[j]} / (y_{E,t}^{[j]} + y_{N,t}^{[j]} + y_{M,t}^{[j]} + y_{U,t}^{[j]})$ denote the percentage of individuals in rotation group j sampled in month t with measured status $X = E, N, M$, or U ; thus $e_t^{[j]} + n_t^{[j]} + m_t^{[j]} + u_t^{[j]}$ exactly equals 100 for every j and every t . Consider an 8-variable panel regression with time fixed effects where the dependent variable is $n_t^{[j]}$, $j = 1, \dots, 8, t = 1, \dots, T$:

$$n_t^{[j]} = \alpha_{nt} + \delta_n j + \alpha_{n1} d_{1t} + \alpha_{n5} d_{5t} + \varepsilon_{nt}^{[j]}. \quad (13)$$

¹⁴For example, the entry in the first row and column is $T^{-1} \sum_{t=1}^T y_{E,t}^{[1]}$.

Here α_{nt} is the time fixed effect for month t , δ_n captures a linear trend across rotations (with increased fraction of N in later rotations captured by $\delta_n > 0$), $d_{1t} = 1$ if $j = 1$ and 0 otherwise allows for something special about the first rotation group, while $d_{5t} = 1$ if $j = 5$ serves a similar function for rotation 5. The fitted value of this regression (with fixed effect $a_{nt} = 0$) is plotted as the red curve in Figure 5. These coefficients capture the tendency for the percentage of individuals classified as N to increase sharply across rotation groups.

Coefficients for panel regressions in which $e_t^{[1]}, \dots, e_t^{[8]}$ are the 8 dependent variables are plotted as the thick black curve in Figure 5. Coefficients when unemployment is the dependent variable are plotted as the dashed yellow line. The rising trend across rotations in N ($\delta_N = 0.0010$) is accounted for by falling trends in E and U ($\delta_E + \delta_U = -0.0012$). The bulges in M in rotation 1 ($\alpha_{M1} = 0.0168$) and rotation 5 ($\alpha_{M5} = 0.0153$) are accounted for by drops in E and N in those rotations.¹⁵

4 Reconciling the inconsistencies.

This section describes how we reconcile the inconsistencies documented in Section 3.

4.1 Rotation-group bias.

We have seen that a given household can give different answers depending on the number of times the household has previously been interviewed. We interpret this as differences in interview technology: the process by which data are obtained differs across rotations, and the numbers from different rotations mean different things. As a first step we summarize these differences in the form of a counterfactual question: if an individual in rotation j had instead been interviewed using the technology i , how would their answers have differed? In this section we show how to answer this question for $i = 1$ and then find the answer for any i . In the next section we ask, which interview technology i should be used as a baseline summary of the data? We identify several reasons why we prefer to use the answers that people give the first time they are interviewed ($i = 1$).

Summarizing the differences in interview technology. Let $\pi_t^{[j]} = (\pi_{E,t}^{[j]}, \pi_{N,t}^{[j]}, \pi_{M,t}^{[j]}, \pi_{U,t}^{[j]})'$ denote

¹⁵These findings are consistent with Krueger, Mas, and Niu's (2017) finding that rotation-group bias is associated with nonresponses and with Bailer's (1975) conclusion that the rotation-group bias of the unemployment rate can be explained by the participation margin.

the observed fraction of individuals who reported status X when interviewed in rotation j in month t . For each $j \in J = \{2, 3, 4\} \cup \{6, 7, 8\}$, of the individuals who reported status X_1 in rotation $j - 1$ in month $t - 1$, some fraction $\pi_{X_1, X_2, t}^j$ are observed to report status X_2 in rotation j for $X_i \in \{E, N, U, M\}$; thus $\pi_{XE, t}^{[j]} + \pi_{XN, t}^{[j]} + \pi_{XU, t}^{[j]} + \pi_{XM, t}^{[j]} = 1$ for all X, t and $j \in J$. Collect these observed probabilities in a matrix

$$\Pi_t^{[j]} = \begin{bmatrix} \pi_{EE, t}^{[j]} & \pi_{NE, t}^{[j]} & \pi_{ME, t}^{[j]} & \pi_{UE, t}^{[j]} \\ \pi_{EN, t}^{[j]} & \pi_{NN, t}^{[j]} & \pi_{MN, t}^{[j]} & \pi_{UN, t}^{[j]} \\ \pi_{EM, t}^{[j]} & \pi_{NM, t}^{[j]} & \pi_{MM, t}^{[j]} & \pi_{UM, t}^{[j]} \\ \pi_{EU, t}^{[j]} & \pi_{NU, t}^{[j]} & \pi_{MU, t}^{[j]} & \pi_{UU, t}^{[j]} \end{bmatrix} \quad j \in J.$$

Notice that each column of $\Pi_t^{[j]}$ sums to unity. For example, for the first column, if someone reported status E when interviewed in rotation $j - 1$, they must have had one of the statuses E, N, M , or U in rotation j . Our constructed data set exactly satisfies the accounting identity

$$\pi_t^{[j]} = \Pi_t^{[j]} \pi_{t-1}^{[j-1]} \quad \text{for all } t \text{ and } j \in J. \quad (14)$$

For an individual who reported status $X^{[j]}$ in rotation j in month t , consider the counterfactual answer that individual would have given if interviewed using the interview technology that was used for rotation 1:

$$r_{X^{[j]}, X^{[1]}, t}^{[j]} = \text{Prob}(\text{would have answered } X^{[1]} \text{ using technology 1 given answered } X^{[j]} \text{ using technology } j).$$

Collect these counterfactual probabilities in a matrix

$$R_t^{[j]} = \begin{bmatrix} r_{EE, t}^{[j]} & r_{NE, t}^{[j]} & r_{ME, t}^{[j]} & r_{UE, t}^{[j]} \\ r_{EN, t}^{[j]} & r_{NN, t}^{[j]} & r_{MN, t}^{[j]} & r_{UN, t}^{[j]} \\ r_{EM, t}^{[j]} & r_{NM, t}^{[j]} & r_{MM, t}^{[j]} & r_{UM, t}^{[j]} \\ r_{EU, t}^{[j]} & r_{NU, t}^{[j]} & r_{MU, t}^{[j]} & r_{UU, t}^{[j]} \end{bmatrix} \quad j \in J.$$

Notice that each column of $R_t^{[j]}$ sums to unity. For example, for the first column, given that an individual reported status E when interviewed in rotation j , they would have to have given one

of the answers E, N, M, U if interviewed using the technology of rotation 1. We can construct matrices $R_t^{[j]}$ that satisfy the condition¹⁶

$$R_t^{[j]} \pi_t^{[j]} = \pi_t^{[1]} \quad \text{for } t = 1, \dots, T \text{ and } j = 2, \dots, 8. \quad (15)$$

From the analysis above, for $j > 1$ we expect $r_{NU,t}^{[j]} > 0$; some of the individuals who report labor status N in rotation j would have reported status U if they had been interviewed for the first time. We also expect $r_{EM,t}^{[j]} > 0$ and $r_{NM,t}^{[j]} > 0$; some of the individuals who were reported as status E or N in rotation j would have been missing using the interview technology of rotation 1.

One can parameterize a matrix $R_t^{[j]}$ that exactly satisfies (15) in an infinite number of ways. In our monthly empirical estimates below we will take the view that rotation bias evolves slowly over time, leading us to replace $R_t^{[j]}$ with an estimate $\bar{R}_t^{[j]}$ where $\bar{R}_t^{[j]}$ does not differ too much from $\bar{R}_t^{[j-1]}$. In this case, $\bar{R}_t^{[j]} \pi_t^{[j]}$ will be close to but not exactly equal to $\pi_t^{[1]}$. In anticipation of this plan, we now parameterize the unrestricted matrix $R_t^{[j]}$ in a way that focuses on what we believe to be the most important features of rotation bias. In Figure 5 we saw that the decline in U across rotations is balanced by a corresponding trend up in N and that differences in M in rotations 1 and 5 correspond to matching drops in E and N . We therefore propose to capture the key differences in interview technology in month t using three parameters $\theta_t^{[j]} = (\theta_{EM,t}^{[j]}, \theta_{NM,t}^{[j]}, \theta_{NU}^{[j]})'$:¹⁷

$$R_t^{[j]} = \begin{bmatrix} 1 - \theta_{EM,t}^{[j]} & 0 & 0 & 0 \\ 0 & 1 - \theta_{NM,t}^{[j]} - \theta_{NU,t}^{[j]} & 0 & 0 \\ \theta_{EM,t}^{[j]} & \theta_{NM,t}^{[j]} & 1 & 0 \\ 0 & \theta_{NU,t}^{[j]} & 0 & 1 \end{bmatrix}. \quad (16)$$

¹⁶For example, the first row states

$$r_{EE,t}^{[j]} \pi_{E,t}^{[j]} + r_{NE,t}^{[j]} \pi_{N,t}^{[j]} + r_{ME,t}^{[j]} \pi_{M,t}^{[j]} + r_{UE,t}^{[j]} \pi_{U,t}^{[j]} = \pi_{E,t}^{[1]}.$$

This equation states that the fraction who reported E in rotation 1 can be viewed as the fraction who reported $X^{[j]}$ in rotation j times the probability someone reporting $X^{[j]}$ would have reported E using technology 1, added across the four possible $X^{[j]}$.

¹⁷We take the (3,3) and (4,4) elements of $R_t^{[j]}$ to be unity because a higher fraction of the population is M or U in rotation 1 than in other rotations. For example, the third equation in (15) states that the fraction missing in rotation 1 is the fraction missing in rotation j plus some portions $\theta_{EM,t}^{[j]}$ and $\theta_{NM,t}^{[j]}$ of the fractions that are E and N in rotation j : $\pi_{M,t}^{[1]} = \pi_{M,t}^{[j]} + \theta_{EM,t}^{[j]} \pi_{E,t}^{[j]} + \theta_{NM,t}^{[j]} \pi_{N,t}^{[j]}$. Note that the normalization of the third and fourth columns of $R_t^{[j]}$ still allows equation (15) to fit exactly the observed average values of every element of $\pi_t^{[j]}$ for every j and t .

The value of $\theta_t^{[j]}$ that causes (15) to hold exactly for every j is given by¹⁸

$$1 - \theta_{EM,t}^{[j]} = \pi_{E,t}^{[1]} / \pi_{E,t}^{[j]} \quad (17)$$

$$\theta_{NU,t}^{[j]} = (\pi_{U,t}^{[1]} - \pi_{U,t}^{[j]} / \pi_{N,t}^{[j]}) / \pi_{N,t}^{[j]} \quad (18)$$

$$1 - \theta_{NM,t}^{[j]} - \theta_{NU,t}^{[j]} = \pi_{N,t}^{[1]} / \pi_{N,t}^{[j]}. \quad (19)$$

Next consider a vector π_t^* that represents the fractions that would have been reported if everyone in month t had been interviewed using technology 1. One reasonable estimate of π_t^* would be $\pi_t^{[1]}$. However, our proposal below will be to estimate π_t^* using all eight of the observed $\pi_t^{[j]}$ under the assumption that the rotation bias parameters in $R_t^{[j]}$ do not change much over time. That is, we will use $\pi_t^{[2]}$ to help improve our estimate of π_t^* by adjusting $\pi_t^{[2]}$ based on the average relation between $\pi_t^{[2]}$ and $\pi_t^{[1]}$ in recent years.

We also construct a counterfactual matrix of transition probabilities Π_t^* that summarize what transitions between labor-force status would have been if all individuals could have been interviewed in both $t-1$ and t using interview technology 1. We require this estimate to satisfy the accounting identity

$$\Pi_t^* \pi_{t-1}^* = \pi_t^*. \quad (20)$$

We have reasonable estimates of π_{t-1}^* and π_t^* (e.g., $\pi_{t-1}^{[1]}$ and $\pi_t^{[1]}$, respectively). The goal is to use the observed transition probabilities $\Pi_t^{[j]}$ and the representation of rotation bias in (15) to construct an estimate of Π_t^* satisfying (20).

Premultiply (14) for $j = 2$ by $R_t^{[2]}$ and use result (15):

$$R_t^{[2]} \Pi_t^{[2]} \pi_{t-1}^{[1]} = R_t^{[2]} \pi_t^{[2]} = \pi_t^{[1]}.$$

¹⁸These equations come from solving rows 1, 2 and 4 of (15). One can show that equations (17)-(19) imply that row 3 of (15) also holds. Add rows 1, 2, and 4 of (15) together to deduce

$$\pi_{E,t}^{[j]} + \pi_{U,t}^{[j]} + \pi_{N,t}^{[j]} - \theta_{EM,t}^{[j]} \pi_{E,t}^{[j]} - \theta_{NM,t}^{[j]} \pi_{N,t}^{[j]} = \pi_{E,t}^{[1]} + \pi_{U,t}^{[1]} + \pi_{N,t}^{[1]}.$$

Subtracting both sides from 1 gives

$$\pi_{M,t}^{[j]} + \theta_{EM,t}^{[j]} \pi_{E,t}^{[j]} + \theta_{NM,t}^{[j]} \pi_{N,t}^{[j]} = \pi_{M,t}^{[1]}$$

as required by the third row of (15). In general, since each column of $R_t^{[j]}$ sums to unity, if elements of $\pi_t^{[j]}$ sum to unity, then the elements of $R_t^{[j]} \pi_t^{[j]}$ also sum to unity: $1' R_t^{[j]} \pi_t = 1' \pi_t = 1$ for 1 a vector of four ones.

In other words, $R_t^{[2]}\Pi_t^{[2]}$ offers one estimate of the counterfactual transition matrix Π_t^* if people could somehow have been interviewed with the same technology in rotation 2 as in rotation 1. It satisfies the internal consistency requirement (20), namely, $\Pi_t^*\pi_{t-1}^{[1]} = \pi_t^{[1]}$ for $\Pi_t^* = R_t^{[2]}\Pi_t^{[2]}$. More generally, premultiplying (14) by $R_t^{[j]}$ we see

$$R_t^{[j]}\Pi_t^{[j]}(R_{t-1}^{[j-1]})^{-1}R_{t-1}^{[j-1]}\pi_{t-1}^{[j-1]} = R_t^{[j]}\pi_t^{[j]}$$

$$R_t^{[j]}\Pi_t^{[j]}(R_{t-1}^{[j-1]})^{-1}\pi_{t-1}^{[1]} = \pi_t^{[1]} \quad \text{for } j \in J \quad (21)$$

where $R_t^{[1]}$ is defined to be the identity matrix. Thus $\Pi_t^{*[j]} = R_t^{[j]}\Pi_t^{[j]}(R_t^{[j-1]})^{-1}$ gives us another estimate satisfying the internal consistency requirement $\Pi_t^{*[j]}\pi_{t-1}^{[1]} = \pi_t^{[1]}$. Our approach will be to estimate Π_t^* so as to be as close as possible to the various estimates $\Pi_t^{*[j]}$ while satisfying all the necessary accounting identities, as described in detail below.

Rotation-bias correction to the full-sample averages. To calculate a full-sample analog to the date t estimate just described, we replace $\pi_t^{[j]}$ in equations (17)-(19) with $\pi^{[j]}$, the average fractions in rotation j across our full sample. This produces the estimates of $\theta^{[j]}$ reported in Table 5. The first row shows that 1-2% of the individuals who get counted as employed in rotations 2-4 or 6-8 would have been missing from the survey if the rotation 1 interview technology had been used. On the other hand, rotation 5 (which follows an 8-month break) reports similar numbers of E as rotation 1 ($\theta_{EM}^{[5]}$ near 0).¹⁹ The second row captures a rising tendency for those who would have been counted as N in later rotations to have been counted as U in the first interview. The third row indicates that a large and rising fraction of those counted N in later rotations would have been M in rotation 1.

Let $\Pi^{[j]}$ be the observed full-sample average transition probabilities into rotation j and $\bar{R}^{[j]}$ be the value obtained by plugging the parameter values in Table 5 into expression (16). We then chose values for the $(n \times n)$ matrix Π^* by minimizing the sum of squared elements of

$$\Pi^{[j]} - (\bar{R}^{[j]})^{-1}\Pi^*\bar{R}^{[j-1]} \quad \text{for } j \in J \quad (22)$$

¹⁹The estimate of $\theta_{EM}^{[5]}$ from equation (17) is actually very slightly negative (-0.0049). The value reported in Table 5 and used in the calculations below sets $\theta_{EM}^{[5]} = 0$. This makes essentially no difference for any results.

$$\pi^{[1]} - \pi^* \quad (23)$$

$$\pi^{[5]} - (\bar{R}^{[5]})^{-1}\pi^* \quad (24)$$

subject to the constraints that all elements of Π^* lie between 0 and 1, each column of Π^* sums to 1, and that π^* is the vector of ergodic probabilities implied by Π^* .²⁰ The resulting estimates of π^* and Π^* are reported in Table 6.

This framework predicts that the fraction of individuals reporting status E, N, M , or U when interviewed using technology j would be given by

$$\hat{\pi}^{[j]} = (\bar{R}^{[j]})^{-1}\pi^*. \quad (25)$$

These predicted shares are compared with the actual shares reported for each rotation in Figure 6. Our representation fits the values in each $\pi^{[j]}$ essentially perfectly.

Our approach also implies a predicted value for the observed fraction of individuals with measured transitions from $X^{[j-1]}$ to $X^{[j]}$:

$$\hat{\Pi}^{[j]} = (\bar{R}^{[j]})^{-1}\Pi^*\bar{R}^{[j-1]}. \quad (26)$$

Figure 7 plots these predicted values along with the actual reported fractions for $j \in J$.²¹ These show a reasonable fit, though not perfect. One could try to model in more detail features such as the tendency for those missing in rotation 1 to be reported as employed in rotation 2 and for those not in the labor force in rotation 1 to be missing in rotation 2. Notwithstanding, our simple parsimonious framework does a reasonable job of capturing transitions.

We defined the value of π^* in terms of the rotation 1 technology. But now that we've found π^* , we can also calculate the answer using any other technology. For example, $(\bar{R}^{[5]})^{-1}\pi^*$ gives

²⁰That is, we minimized the sum of squares of the $96 = 16 \times 6$ elements in (22) plus the sum of squares of the 8 elements in (23) and (24). The vector π^* is also a function of Π^* using expression [22.2.26] in Hamilton (1994):

$$B = \begin{bmatrix} I_4 - \Pi^* \\ 1' \end{bmatrix}$$

$$\pi^* = (B'B)^{-1}B'e_5$$

where $1'$ denotes a (1×4) vector of ones and e_5 denotes column 5 of I_5 .

²¹Note we do not offer a predicted value for transitions from $X^{[4]}$ to $X^{[5]}$ since there are 8 intervening months between rotations 4 and 5.

the answer in terms of the rotation 5 technology. The BLS approach, which simply averages the rotations together, is implicitly reporting results in terms of an “average” technology, which in our formulation would be described as $\pi^{**} = \tilde{R}^{-1}\pi^*$ for $\tilde{R}^{-1} = (1/8)\sum_{j=1}^8(\bar{R}^{[j]})^{-1}$. Appendix Table A-5 reports π^{**} and $\Pi^{**} = \tilde{R}^{-1}\Pi^*\tilde{R}$, our estimates of the full-sample averages and transition probabilities if all individuals had been surveyed using the average interview technology.

Month-by-month corrections for rotation bias. For applying this approach to monthly data, we take the view that the rotation bias parameters evolve slowly over time, implemented using the principle of exponential smoothing. Our first step is to construct weighted moving averages of the counts of individuals in each labor-force status in each rotation,

$$\bar{y}_{X,t}^{[j]} = (1 - \lambda)y_{X,t}^{[j]} + \lambda\bar{y}_{X,t-1}^{[j]},$$

where $y_{X,t}^{[j]}$ denotes the observed weighted number of individuals reporting labor status $X \in \{E, N, M, U\}$ in rotation j in month t . For $\lambda = 1$, this method would reproduce the full-sample averages just reported. For $\lambda = 0$, it would amount to estimating values for each month in isolation of all the others. We set $\lambda = 0.98$, which means that observations 3 years prior to t receive half the weight of observation t in determining the smoothed count $\bar{y}_{X,t}^{[j]}$.²² We then calculated the corresponding smoothed fractions as

$$\bar{\pi}_{X,t}^{[j]} = \bar{y}_{X,t}^{[j]} / \left(\bar{y}_{E,t}^{[j]} + \bar{y}_{N,t}^{[j]} + \bar{y}_{M,t}^{[j]} + \bar{y}_{U,t}^{[j]} \right).$$

From these we calculated time-varying rotation-bias parameters as

$$\theta_{EM,t}^{[j]} = \max \left\{ 1 - \left(\bar{\pi}_{E,t}^{[1]} / \bar{\pi}_{E,t}^{[j]} \right), 0 \right\}$$

$$\theta_{NU,t}^{[j]} = \max \left\{ \left(\bar{\pi}_{U,t}^{[1]} - \bar{\pi}_{U,t}^{[j]} \right) / \bar{\pi}_{N,t}^{[j]}, 0 \right\}$$

$$\theta_{NM,t}^{[j]} = \max \left\{ 1 - \theta_{NU,t}^{[j]} - \left(\bar{\pi}_{N,t}^{[1]} / \bar{\pi}_{N,t}^{[j]} \right), 0 \right\},$$

and exponentially smoothed these as well. For example, $\bar{\theta}_{EM,t}^{[j]} = (1 - \lambda)\theta_{EM,t}^{[j]} + \lambda\bar{\theta}_{EM,t-1}^{[j]}$.

²²That is, $0.98^{36} = 0.48$. We started the recursion by setting $\bar{y}_{X,1}^{[j]} = (1/36)\sum_{t=1}^{36} y_{X,t}^{[j]}$ the average of the first three years of observations.

The resulting series for $\bar{\theta}_{EM,t}^{[j]}$, $\bar{\theta}_{NU,t}^{[j]}$, and $\bar{\theta}_{NM,t}^{[j]}$ are plotted in Figure 8. The value of $\bar{\theta}_{EM,t}^{[j]}$, which characterizes the tendency to record people as E in rotation j who would have been M in rotation 1, has fallen somewhat over time. By contrast, $\bar{\theta}_{NU,t}^{[j]}$, which governs the tendency of people who would have been counted as U in earlier rotations to be designated as N in later rotations, has increased over time. The third parameter, $\bar{\theta}_{NM,t}^{[j]}$, which characterizes the tendency of someone who would have been counted as M in rotation 1 to be counted as N in later rotations, has not changed much over time.

Plugging the values for $\bar{\theta}_{EM,t}^{[j]}$, $\bar{\theta}_{NU,t}^{[j]}$, and $\bar{\theta}_{NM,t}^{[j]}$ into (16) gives a value of $\bar{R}_t^{[j]}$ for each j and t . Our procedure was to proceed iteratively through the data, choosing Π_t^* for each t to minimize the errors in the following equations:

$$\Pi_t^{[j]} - (\bar{R}_t^{[j]})^{-1} \Pi_t^* \bar{R}_{t-1}^{[j-1]} \quad \text{for } j \in J = \{2, 3, 4\} \cup \{6, 7, 8\} \quad (27)$$

$$\pi_t^{[1]} - \Pi_t^* \pi_{t-1}^* \quad (28)$$

$$\pi_t^{[5]} - (\bar{R}_t^{[5]})^{-1} \Pi_t^* \pi_{t-1}^*. \quad (29)$$

We set the initial value of π_t^* for observation $t = 1$ as $\pi_1^* = \pi_1^{[1]}$. For each $t = 2, 3, \dots$ we choose the 16 elements of Π_t^* so as to minimize the sum of squares of the 104 terms in (27)-(29) subject to the constraints that each element of Π_t^* is between 0 and 1 and each column of Π_t^* sums to 1. Given Π_t^* we then calculated

$$\pi_t^* = \Pi_t^* \pi_{t-1}^*$$

and proceeded to the next observation $t + 1$. The resulting time series $\{\pi_t^*\}_{t=1}^T$ gives the estimate for month t of the fractions that would have reported each status using interview technology 1, and is the starting point for the adjusted estimates described below.

4.2 Choosing a baseline interview technology.

The framework in Section 4.1 allows us to reconcile stocks and flows in the CPS data and summarize that reconciliation using any interview technology. In practice we need to choose a

particular technology as a baseline. In this section we review the reasons why we recommend using the first-interview definition of labor-force status.

Disengagement. The tendency to report a higher incidence of unemployment the first time people are asked has also been observed in the Netherlands (van den Brakel and Krieg, 2015) and New Zealand (Silverstone and Bell, 2010). One possible explanation is that people become less engaged the more times they are interviewed and tend toward answers that they think will end the interview more quickly. For example, the CPS interview is more onerous if the respondent says that they have worked at more than one job. The number of people reporting more than one job drops sharply across rotations (Halpern-Manners and Warren, 2012; Hirsch and Winters, 2016). The CPS questionnaire also routes people over age 50 who say they are retired through an abbreviated set of labor-force questions.²³ It is interesting to note that more than all of the increasing incidence of N in later rotations can be explained by larger numbers of people who say they are retired or disabled. Specifically, note from row 1 in Table 7 that the average fraction of the population categorized as N is 1.3% higher in rotations 2-8 than in rotation 1. Row 2 shows that the fraction of the population categorized as N and retired is 1.0% higher in 2-8, and row 3 shows that the fraction categorized as N and disabled is 0.6% higher.

This raises the possibility that some of the people who had reported U in rotation 1 hoped to end the interview more quickly in later interviews if they claimed to be retired or disabled. We can observe in the data that those who are allegedly retired or disabled in rotations 2-7 are more likely to return to the labor force (that is, to report E or U) the following month than are the retired or disabled in rotation 1 (see row 4 of Table 7). This observation is consistent with the inference that some of the additional individuals in later rotations who are designated as N are in an objective sense still in the labor force.

Stigma. Another possible explanation suggested by Halpern-Manners and Warren (2012) is that some people may perceive a stigma in reporting to an official government agency that they are continually searching for a job without success. This of course could interact with the disengagement effect – someone who feels some stigma associated with their status may become less engaged with the interview process than the general population. This could lead some respondents to report in subsequent interviews that they did not actively search for work even though they did, which would

²³Current Population Survey Interviewing Manual, April 2015, p. B3-3.

show up as an increase in N and decrease in U in later rotations. The CPS allows one member of the household to report the labor-force status for all the adults living there. It is noteworthy that unemployment falls much more quickly across rotations among individuals who are reporting their own status compared to individuals whose status is reported by a proxy, consistent with Halpern-Manners and Warren’s hypothesis. Self-responders account for half of the total observations but two-thirds of rotation bias (see rows 6 and 7 of Table 7).

Demographic evidence. An alternative hypothesis is that some individuals are confused by the questions or learn the meaning better as interviews are repeated. Some interesting evidence on this comes from differences in rotation bias across demographic groups and over time. Following Krueger, Mas and Niu (2017), we summarize the magnitude of rotation bias by the slope of a regression of the unemployment rate for rotation j on a constant and the month-in-sample j . For the subsample 2001:7 to 2008:6, the slope of this regression is -0.11 , meaning that on average, rotation j reported an unemployment rate that is 0.11 percentage points lower than that reported by rotation $j - 1$. For those without a high school degree the slope was -0.21 while for those with a college degree the slope was only -0.05 . These differences might seem to lend support to the confusion/learning hypothesis. Nonetheless, we note that even for those with college degrees, the slope is still highly statistically significant, with a standard error below 0.01.

Another feature that distinguishes less educated workers is that they have higher unemployment rates than the general population. The top panel of Figure 9 illustrates the importance of this graphically. The horizontal axis plots the average unemployment rate as measured by BLS for a particular demographic group and sample period. The vertical axis plots the absolute value of the slope from the Krueger, Mas and Niu regression. For the overall population, the average unemployment rate over 2001-2008 was 5.3%, represented by the point A1 = (5.3, 0.11) in the figure. College graduates (represented by the point C1) had an average unemployment rate over this period of 2.6%, about half that of the overall population, just as the slope for this group is about half that for the overall population for this period. The average unemployment rate for those without a high school diploma, 11.2%, is about twice that for the overall population, as is the slope for this group, as represented by point L1 in the figure.

The slope is different across demographic groups and also changes over time. Consider a second sample, 2008:7-2014:6, that includes the high unemployment rates in the aftermath of the Great

Recession. For almost every demographic group, the unemployment rate is higher in the second subsample than in the first, and the magnitude of rotation bias goes up by a roughly proportionate amount. It seems implausible that the college-educated (S2 and C2) were more confused by the questions during the Great Recession than they had been in the earlier decade. Nor is the increase in rotation bias part of a long-term trend. Over the 2014:7-2020:1 subsample, we see slopes and unemployment rates come back down together for every group.

We can summarize this regularity with an OLS regression fit to the 27 data points in the top panel of Figure 9 (standard errors in parentheses):

$$s_{iT} = \underset{(0.010)}{0.029} + \underset{(0.0013)}{0.0126}u_{iT} + \hat{\varepsilon}_{iT} \quad R^2 = 0.80. \quad (30)$$

Here s_{iT} is the negative of the Krueger-Mas-Niu slope coefficient²⁴ for demographic group i and subsample T and u_{iT} is the average unemployment rate for that group. Most of the variation in the slope across demographic groups and across time can be explained by differences in the unemployment rate across groups and across time. The data are suggestive of a universal law: if a percentage u_{it} of a group i at date t are truly unemployed, 1.26% of those unemployed individuals will no longer be classified as unemployed in each subsequent interview. This law seems to hold for every education group, demographic group, and point in time. Rotation bias can be explained statistically solely on the basis of the underlying true unemployment rate and has nothing to do with education or age. This observation casts doubt on the hypothesis that rotation bias arises from confusion or learning about the questions and is exactly what we would expect to see if stigma and disengagement are the main explanation for rotation bias.

Rotation bias in the employment to population ratio. We observe differences across rotations not just in the unemployment rate but also in the employment rate. The employment-to-population ratio is 0.9 percentage points lower on average for rotations 2-8 compared to rotation 1 (see row 8 of Table 7). Denote by \dot{E}_t the subset of individuals who worked only part-time in month t or who are usually employed but didn't work in t . Row 9 shows that rotation bias in E can be entirely explained by the rotation bias in \dot{E} . This is another labor-force status for which we might expect to see some stigma and less engagement as people are re-interviewed.

²⁴That is, $s_{iT} = -\hat{\beta}_{iT}$ in the regression $u_{it}^{[j]} = \alpha_{iT} + \beta_{iT} \cdot j + \varepsilon_{it}$ for $t \in T$.

The bottom panel of Figure 9 reproduces the analysis in the top panel with U now replaced by \dot{E} . Here $\dot{e}_{it}^{[j]}$ is the fraction of people in rotation j , demographic group i and month t who only work part time in t or report that they are usually employed but not in month t . Let \dot{s}_{iT} be the negative of the slope coefficient from a regression of $\dot{e}_{it}^{[j]}$ on a constant and j over subsample T . The bottom panel of Figure 9 shows the scatterplot relating \dot{s}_{iT} to the average value of \dot{e}_{it} for that group and subsample. A regression line summarizing the relationship is

$$\dot{s}_{iT} = \underset{(0.035)}{0.014} + \underset{(0.0021)}{0.0077}\dot{e}_{iT} + \hat{\varepsilon}_{iT} \quad R^2 = 0.34. \quad (31)$$

This does not have as good a fit as (30) – demographics and sample period play a more important role in rotation bias in the unemployment rate than in the employment-to-population ratio. Notwithstanding, the slope coefficient in (31) is highly statistically significant.

College-educated individuals are significant outliers in the second panel of Figure 9. Part-time work for these individuals may involve consulting or programming for which they may feel little or no stigma compared to less-educated individuals whose part-time work may be cleaning or serving. In fact, the slope coefficients corresponding to the height of C1, C2 and C3 are each far from statistically significant. When these three observations are dropped, the R^2 rises to 0.49.

Evidence from reported durations of job search. One of the main objectives of our study is to reconcile the discrepancies between different CPS statistics. One glaring inconsistency is that reported durations of unemployment are much longer than could be consistent with observed probabilities of UU continuations. Because more U get counted as N as we increase the number of interviews j , if we were to reconcile stocks and flows on the basis of the interview $j - 1$ technology, some of the observed UN transitions between rotation $j - 1$ and j would be interpreted as UU continuations, decreasing the inconsistency between UU continuations and reported durations. By contrast, if we were to standardize on the basis of interview j technology, some of the reported UN transitions would be interpreted as NN continuations, increasing the inconsistency between reported durations and the probability of a UU continuation. Normalizing on the basis of any interview technology $j > 1$ would reduce the number of imputed UU continuations and thus increase the discrepancy between UU continuation probabilities and reported durations.²⁵ Using

²⁵For example, the first-interview measure implies an unemployment-continuation probability of $\pi_{UU}^*/(1 - \pi_{UM}^*) = 55.7\%$ after correcting for rotation bias. By contrast, if we were to use rotation-bias-corrected transition probabilities

the first-interview definition of unemployment helps resolve the inconsistency between reported durations and observed UU continuations relative to a standardization based on any other interview technology.

The role of missing observations. Some have conjectured that rotation bias might arise from unemployed individuals exiting the sample more quickly than others. But rows 10-16 of Table 7 establish that the effect is instead directly related to the number of times the household has been interviewed. There are some people who were interviewed for the first time when the address would have been in rotation 2, for example because the individual moved into the address. The unemployment rate for these individuals is reported in row 11. Others were missing in both 1 and 2 and are being asked the questions for the first time in rotation 3 (row 12). For every group, we see the highest unemployment rate the first time people are asked the questions and a drop across each follow-up interview.²⁶

Moreover, we observe an *increase* across rotations in the total number of individuals who are designated as not in the labor force (row 17 of Table 7). This cannot be people dropping out of the survey, but must come from some people changing their answers. Yet another way to get at this question is to look at the subset of individuals who gave answers in both rotation 1 and rotation 2. Row 1 of Table 8 shows that the unemployment rate for this group was 6.50% the first time they were asked the question and 6.23% the second time. Row 2 shows that among individuals who were sampled in both rotation 2 and rotation 3, the unemployment rate was 6.22% in rotation 2 and 5.99% in rotation 3.²⁷ The same pattern of the reported unemployment rate to drop among a fixed group of individuals whenever the household is asked the same questions a second time is seen in each of the subsequent rows of Table 8 as well.

Reconciliation with Krueger, Mas and Niu. The evidence in Krueger, Mas and Niu (2017) is sometimes interpreted as showing that rotation bias does not result from individuals being asked

π_{X_1, X_2}^{**} in Table A-5 based on the average interview technology, we would calculate an implied unemployment-continuation probability of $\pi_{UU}^{**}/(1 - \pi_{UM}^{**}) = 54.9\%$.

²⁶Indeed, the reported unemployment rate among people being asked the questions for the first time when in rotation 2 (7.4%) is even higher than the unemployment rate among people being asked the questions the first time when in rotation 1 (6.6%). This is a consequence of the fact that individuals who are M in some month of the survey are more likely than the general population to be U in the months when they are sampled.

²⁷The average unemployment rate in row 2, column 2 of Table 8 (6.22%) is not quite the same as in row 3, column 1 (6.23%) because the set of individuals who were neither $M2$ nor $M3$ (which is the set of people who are counted in row 2) is not quite the same as the set of individuals who were neither $M3$ nor $M4$ (which is the set of individuals who are counted in row 3).

the same question multiple times. Krueger, Mas and Niu interpreted the duration of job search as a measure of the number of times an individual had previously reported being unemployed. But duration of job search is not a reliable indicator of the number of times people have answered the questions in earlier rotations. Of people in our sample who responded in both rotations 1 and 2, 30% of the U individuals in rotation 2 who reported unemployment durations 9 weeks or longer had been counted as E or N in rotation 1 (4 weeks earlier). Krueger, Mas and Niu found that the biggest difference between rotations 1 and 2 comes from comparing people who report being unemployed with a duration less than 5 weeks ($U^{1.4}$) in rotation 1 with people who are $U^{1.4}$ in rotation 2. This is not an apples-to-apples comparison. In our 2001-2018 sample, the durations in rotation 1 are all solicited explicitly, whereas the durations for UU continuations into rotation 2 are imputed to be a number greater than 4 weeks. Thus by construction no one who is $U^{1.4}$ in rotation 2 could have been unemployed in rotation 1. Any statistic that conditions on not being U the previous month is selecting a subset of individuals who have a lower unemployment rate than the general population, which explains why $U^{1.4}$ in rotation 2 would be expected to be a smaller number than $U^{1.4}$ in rotation 1. Our data set contains a total of 39,000 individuals who were $U^{1.4}$ in rotation 1 but only 30,000 who were $U^{1.4}$ in rotation 2. By contrast, we have 28,000 $U^{5.14}$ in rotation 1 and 34,000 in rotation 2. This suggests that most of the “missing” $U^{1.4}$ in rotation 2 are being classified as $U^{5.14}$ on the basis of the BLS duration imputation but would have reported $U^{1.4}$ if allowed. The same pattern is seen in comparing rotations 5 and 6. Two-thirds of the drop in $U^{1.4}$ between 5 and 6 is accounted for by the rise in $U^{5.14}$.

Before 1994, durations for all individuals (including UU continuations) were directly solicited rather than imputed. A striking finding in Krueger, Mas and Niu’s Figure 4 is that rotation bias among $U^{1.4}$ individuals was virtually nonexistent prior to 1994 and then appeared suddenly and dramatically when the BLS began imputing durations to UU continuations in 1994. Their figure shows that this break also coincides with a *decrease* in rotation bias in 1994 for $U^{5.14}$ individuals. We conclude that reported and imputed unemployment durations cannot be used in the way suggested by Krueger, Mas and Niu to identify the effects of being asked the survey questions multiple times.²⁸

²⁸Others have suggested that rotation bias might result from a difference between phone interviews and in-person interviews. For example, it is possible that respondents might want to impress the interviewer by showing their effort for job search when jobless, which would overstate the unemployment rate from personal interviews. However, the data suggest to us that this is an unlikely explanation. First, both the first and fifth rotation groups are typically surveyed in person, yet individuals in rotation 5 have significantly lower unemployment rates than those in rotation

Conclusion. We again emphasize that the method we developed in equations (27)-(29) can be used to reconcile stocks with flows on the basis of any interview technology j . The evidence reviewed in this section leads us to recommend the first-interview concept of labor-force status ($j = 1$) as the one that should be used. The most important factor in rotation bias appears to be stigma and disengagement on the part of individuals who are unemployed or employed part time. Evidence in support of this hypothesis includes (1) characteristics of people claiming to be retired or disabled, (2) differences between self-reported answers and answers reported by proxy, (3) observed differences in unemployment rates and part-time employment rates across different demographic and education groups and different sample periods, and (4) comparing answers for missed interviews and matched interviews.

4.3 Reconciling labor-force status with reported unemployment durations.

The inconsistency between reported labor-force status and duration of job search establishes clearly that there must be some errors in either labor-force status or in duration. We argue that it would be inappropriate to try to resolve this inconsistency by completely ignoring duration data. Reported duration signals some important and verifiable information about the individual's true circumstances, as evidenced by the fact that reported duration is a strong statistical predictor of whether the individual will still be unemployed next month (see the top panel of Figure 4). Our final estimates will make corrections to both labor-force status and duration in an effort to resolve the inconsistency. One issue that is key in addressing this inconsistency is the unavoidable gray area in the distinction between someone who is not employed but actively looking for a job (U) versus someone who is truly not in the labor force (N). In this section we review evidence that persuades us that many individuals counted as N by BLS are more accurately interpreted as actively looking for work.

Forecasting evidence about NU^{5+} transitions. Consider individuals who transition from N at $t - 1$ to U at t with a duration of job search reported at t to be 5 weeks or longer, hereafter denoted $N_{t-1}U_t^{5+}$. We first note that the incidence of such transitions can be quite accurately

1 (see column 6 of Table 4). Second, within rotation 5, individuals report significantly lower unemployment rates the more times they have previously been interviewed (see rows 9-12 of column 5 of Table 7). Third, rotation bias was observed during the time when all the interviews were conducted in person (see for example Hansen et al., 1955). For these reasons, we conclude that the mode of interview is unlikely to be the key explanation for rotation bias.

predicted on the basis of the previously observed rate of medium- and long-term unemployment. Let $m_{N,t-1}^\# = 100 \cdot N_{t-1}U_t^{5+}/(E_t + U_t + N_t + M_t)$ be the percentage of individuals who will make an NU^{5+} transition between $t - 1$ and t and let $u_{t-1}^{5+} = 100 \cdot U_{t-1}^{5+}/(E_{t-1} + U_{t-1})$ be the medium- to long-term unemployment rate reported by the BLS in month $t - 1$. The value of u_{t-1}^{5+} turns out to be an excellent predictor of $m_{N,t-1}^\#$, as seen in the following regression (Newey-West standard errors with 12 lags in parentheses):

$$m_{N,t-1}^\# = \underset{(0.0018)}{0.0690} + \underset{(0.0079)}{0.0682}u_{t-1}^{5+} + \hat{\varepsilon}_{t-1} \quad R^2 = 0.91. \quad (32)$$

Whenever one observes a high value of u_{t-1}^{5+} one can quite accurately predict that many of the individuals who are classified as N_{t-1} that month will report in the following month that they have been looking for work for longer than 5 weeks. This observation invites us to ask whether some of the long-term unemployed at date $t - 1$ are being misclassified as N_{t-1} .

We can also examine at the individual level what a person's report of $N_{t-1}U_t^{5+}$ predicts about their own labor-force status at $t + 1$. The first column of Table 9 examines UUU continuations in months $t - 2$, $t - 1$, and t for which the reported durations would be consistent with a true UUU continuation.²⁹ As we go down the rows, the history is consistent with a longer initial duration in month $t - 2$. Our framework would predict that the probability of being employed in month $t + 1$ would decrease as we move down the rows. This is because type 2 individuals, who have a lower probability than type 1 of becoming employed at $t + 1$, make up a larger fraction of the pool at t as we move down the rows.³⁰ This is exactly what we observe in the data. The third column looks at individuals with an intervening N status in month $t - 1$ but with the same U durations in $t - 2$ and t as in column 1. These probabilities also tend to decrease as we move down rows. The job-finding prospects for someone who begins a UUU stretch with reported initial duration of 5 to 14 weeks (16%) is similar to that for somebody who begins a UNU stretch with duration 5-14 weeks (15%), as are the probabilities for someone beginning with more than 26 weeks (8% versus 7%, respectively).

Demographic evidence on NU^{5+} transitions. It is also interesting to examine how $m_{Nt}^\#$ varies

²⁹For example, $U_{t-2}^{1.4}, U_{t-1}^{5.14}, U_t^{5.14}$ refers to someone who reported being newly unemployed in $t - 2$ and being unemployed between 5 and 14 weeks in $t - 1$ and t .

³⁰See Ahn and Hamilton (2020).

across different demographic groups. We summarize this evidence in the top panel of Figure 10. Over the 2001:7-2008:6 subsample, $m_{Nt}^\#$ averaged 0.44% for those without a high school diploma but only 0.15% for college graduates, a fact that could lead some to conclude that better educated individuals are less prone to misreport their time spent searching for a job. But we also observe that the value of $m_{Nt}^\#$ is higher for every group in the 2008:7-2014:6 subsample, correlating very strongly with the higher rates of u_t^{5+} for every group in this subsample. For example, for college graduates the average values of u_t^{5+} and $m_{Nt}^\#$ during the Great Recession were about the same as the average values of these two variables for the overall population prior to the recession. When their unemployment rate is high, college graduates look just like the rest of the population when their unemployment rate is high. Both $u_{cg,t}^{5+}$ and $m_{cg,Nt}^\#$ returned to their historical values when the recession was over. A regression summarizing this relation is

$$m_{i,NT}^\# = \underset{(0.0251)}{0.0629} + \underset{(0.0045)}{0.0609}u_{iT}^{5+} + \hat{\varepsilon}_{iT} \quad R^2 = 0.88. \quad (33)$$

Almost all of the variation in observed NU^{5+} transitions across demographic groups and across time can be explained by the simple hypothesis that when more people are experiencing long-term unemployment, the number of unemployed who are incorrectly classified as not looking for a job is higher.

Note moreover that the coefficients from the demographic regression (33) are remarkably similar to those for the aggregate time-series forecasting relation (32). We illustrate this point visually in the bottom panel of Figure 10. For each month $t - 1$, we would predict on the basis of the demographic regression a value for next month's aggregate $N_{t-1}U_t^{5+}$ transitions to be $0.0629 + (0.0609)u_{i,t-1}^{5+}$. The figure compares this with the actual transitions $m_{N,t-1}^\#$. The predicted $\hat{m}_{N,t-1}^\#$ is on average lower than the actual $m_{N,t-1}^\#$, as a result of the fact that A1-A3 fall a little above the demographic regression line in the top panel. Notwithstanding, the fit is quite remarkable. One can predict $N_{t-1}U_t^{5+}$ transitions quite accurately, for any demographic group or any point in time, solely on the basis of u_{t-1}^{5+} .

Additional evidence on NU^{5+} transitions. There are a number of other reasons to regard individuals who report an $N_{t-1}U_t^{5+}$ transition as having been unemployed at time $t - 1$. First, when asked at time t , "how long have you been looking for work?", their answer (more than 4

weeks) indicates that the individual’s own perception at t is fully consistent with characterizing them as U at $t - 1$. Second, the cross-sectional distribution of reported unemployment durations among those making $N_{t-1}U_t^{5+}$ transitions is remarkably similar to that for those who reported U_{t-1} . Our estimate in Table 1 of p_2 , the key parameter summarizing perceived duration for the long-term unemployed, is 0.9745 for $N_{t-1}U_t^{5+}$ individuals and 0.9737 for U_{t-1} . Further details of the cross-sectional distribution, as summarized both by our parametric model in columns 1 and column 5 in Table 1 and in the raw data in Figures 2 and 3, are strikingly similar. Third, the objective probability of being employed the next period is similar across the two groups: $P(E_{t+1}|N_{t-1}, U_t^{5+}) = 12.5\%$ versus $P(E_{t+1}|U_{t-1}, U_t^{5+}) = 15.5\%$, in sharp contrast for example to $P(E_{t+1}|E_{t-1}, U_t^{5+}) = 37.8\%$. Fourth, information the individuals gave at $t - 1$ would also identify many of the $N_{t-1}U_t^{5+}$ transitions as more attached to the labor force than typical N_{t-1} . Specifically, people who are not in the labor force are asked whether they want a job. Only 5.3% of all N_{t-1} answered this question yes, whereas 44% of $N_{t-1}U_t^{5+}$ answered the question yes at $t - 1$. The indication that a person wants a job (WJ) is furthermore an objective predictor that they will find one. For example, $P(E_{t+1}|N_{t-1}^{WJ}, U_t) = 14.1\%$ versus $P(E_{t+1}|U_{t-1}, U_t) = 15.5\%$.³¹

Based on these considerations, our recommendation is to classify observed $N_{t-1}U_t^{5+}$ transitions as having been U rather than N at $t - 1$. This adjustment is closely related to that recommended by Rothstein (2011), Elsby et al. (2011), Elsby, Hobijn, and Şahin (2015), and Farber and Valletta (2015) who reclassified all UNU as UUU . Our approach differs from theirs in that we utilize the reported duration of unemployment when correcting classification errors. The adjustment just described would only classify UNU as UUU if the final U reports a duration of job search greater than 4 weeks.

Interpreting flows from long-term unemployment into N . If we are correct that some of the people who are currently counted as N are better classified as U , it also means that some UN transitions are really UU continuations. In Section 3.3 we found that the discrepancy between reported unemployment durations and objective unemployment-continuation probabilities mainly

³¹Recently the Federal Reserve Bank of New York has added detailed questions to their Survey of Consumer Expectations about an individual’s search effort, search methods and outcomes, and the incidence of informal recruiting methods. Faberman et al. (2019) find that if one defines unemployment to mean someone who actively searched and is available for work, the unemployment rate in the U.S. over October 2013 to December 2017 would have been 1.7% higher on average than the figures reported by the BLS. This is close to the figure implied by our final adjustment, which is 2.1% higher than the BLS figure over this period. Faberman et al.’s measure does not account for nonrandom missing observations, which could explain the 0.4% difference between their estimate and ours.

comes from $\gamma_{2,UU}$, the objective unemployment-continuation probability for type 2 individuals. Here we explore whether a fraction ξ_{UN} of the $\gamma_{2,UN}$ transitions should be regarded as UU continuations. Since type 2 individuals account for 95% of those unemployed for 15 weeks and over (hereafter, U^{15+}), we look for evidence in the observed outcomes in month t of individuals who were U^{15+} in $t-2$ and N in $t-1$.

Someone with a history $U_{t-2}^{15+}N_{t-1}$ has a 21.4% probability of being U^{15+} in t . We argued above that such an individual, having been observed to be $N_{t-1}U_t^{15+}$, should be classified as U at $t-1$. This means that any $U_{t-2}^{15+}N_{t-1}U_t^{15+}$ sequence is really UUU . Thus at a minimum an average fraction $\xi_{UN} > 0.214$ of $U_{t-2}^{15+}N_{t-1}$ should be regarded as UU continuations.

But $U_{t-2}^{15+}N_{t-1}$ individuals are special not just in their objective probability of returning to unemployment but also in their probability of successfully landing a job. Someone with a $U_{t-2}^{15+}N_{t-1}$ history has a 7.69% probability of being employed at t , far higher than usually observed for individuals classified as N_{t-1} ($P(E_t|N_{t-1}) = 4.10\%$). Suppose we view $U_{t-2}^{15+}N_{t-1}$ individuals as a mixture of two populations, with a fraction ξ_{UN} having the same employment probability in month t as someone who is observed to be $U_{t-2}^{15+}U_{t-1}^{15+}$, and the remainder with the same employment probability as someone who is truly out of the labor force in $t-1$ as represented by a history of $N_{t-2}N_{t-1}$:

$$P(E_t|U_{t-2}^{15+}, N_{t-1}) = \xi_{UN}P(E_t|U_{t-2}^{15+}, U_{t-1}^{15+}) + (1 - \xi_{UN})P(E_{t+1}|N_{t-2}, N_{t-1})$$

$$0.0769 = 0.1094\xi_{UN} + 0.0222(1 - \xi_{UN}).$$

This equation gives an estimate of $\xi_{UN} = 0.627$, which would imply an objective unemployment-continuation probability for type 2 individuals of $\gamma_{2,UU} + \xi_{UN}\gamma_{2,UN}$.

Our proposal is thus to classify some of the N_t as U_t based on a subsequent observation of U_{t+1}^{5+} , and also to classify some of the N_t as U_t based on a previous observation of U_{t-1}^{15+} . We then make a subtraction to correct for double counting $U_{t-1}^{15+}N_tU_{t+1}^{5+}$.

4.4 Nonrandom missing observations.

The conventional approach simply throws out missing observations, which amounts to assuming that those missing from the survey are just like those included. However, our rotation-bias corrected probabilities Π^* in Table 6 show that someone who is employed has a 6.4% probability of being missing in the next month, whereas someone who is unemployed has 8.8% probability. Of those making ME , MN , or MU transitions, 5.8% are unemployed, although the unemployed only comprise 4.3% of the observed E , N , or U on average. In addition, of those making MU transitions, 60% claim that they have been searching for work longer than 4 weeks. In sum, missing individuals are more likely to be unemployed than a typical person in the observed data.

Our M category includes the out-of-scope population, for example, people who leave the sample for reasons such as death, imprisonment, or enlistment in the army. Such individuals would show up in our data set as EM , NM , or UM transitions. Our procedure does not make any adjustment to labor-force measures for such individuals. Instead, our adjustments will be based solely on individuals who were M the previous month and are E , N , or U during the current month. This category does include individuals who were 15 in the previous month but became 16 in the current month, and those who were in the armed force in the previous month but now a civilian. However, we can directly observe these flows from the microdata, and the fractions of these observations are negligible (less than 0.1% of civilian non-institutional population). Hence, it should not affect our estimates significantly.

To correct for the bias coming from nonrandom missing observations, we impute a labor-force status in month $t - 1$ to individuals observed to make ME , MN , or MU transitions into period t . Suppose that some fraction m_E of those missing in month $t - 1$ are just like those who were counted as employed that month in terms of their transition probabilities, while fractions m_N or m_U share the same transition probabilities as those counted as N or U . We regard the remaining $m_M = 1 - m_E - m_N - m_U$ as “dormant observations” in the sense of having zero probability of being recorded as E , N , or U in month t .³² The probabilities of observing ME , MN , and MU

³²This would include people who are in the military, incarcerated, moved away from the address, or yet to move in, for example.

transitions would then be given by

$$\begin{bmatrix} \pi_{ME}^* \\ \pi_{MN}^* \\ \pi_{MU}^* \end{bmatrix} = \begin{bmatrix} \pi_{EE}^* & \pi_{NE}^* & \pi_{UE}^* \\ \pi_{EN}^* & \pi_{NN}^* & \pi_{UN}^* \\ \pi_{EU}^* & \pi_{NU}^* & \pi_{UU}^* \end{bmatrix} \begin{bmatrix} m_E \\ m_N \\ m_U \end{bmatrix}. \quad (34)$$

This system of equations can be solved to find $(m_E, m_N, m_U) = (0.0975, 0.0475, 0.0117)$. Our suggested correction for nonrandom missing observations for the full-sample is then³³

$$\begin{bmatrix} \pi_E^* + \pi_M^* m_E \\ \pi_N^* + \pi_M^* m_N \\ \pi_U^* + \pi_M^* m_U \end{bmatrix}.$$

To obtain monthly estimates, we use Π_t^* to solve for m_{t-1} :

$$\begin{bmatrix} \pi_{ME,t}^* \\ \pi_{MN,t}^* \\ \pi_{MU,t}^* \end{bmatrix} = \begin{bmatrix} \pi_{EE,t}^* & \pi_{NE,t}^* & \pi_{UE,t}^* \\ \pi_{EN,t}^* & \pi_{NN,t}^* & \pi_{UN,t}^* \\ \pi_{EU,t}^* & \pi_{NU,t}^* & \pi_{UU,t}^* \end{bmatrix} \begin{bmatrix} m_{E,t-1} \\ m_{N,t-1} \\ m_{U,t-1} \end{bmatrix}.$$

We also smooth these as

$$\bar{m}_{X,t} = (1 - \lambda)m_{X,t} + \lambda\bar{m}_{X,t-1}.$$

The $m_{X,t}$ parameters have more high-frequency movement than terms like $\theta_{EM,t}$. We accordingly use a shorter effective window by setting $\lambda = 0.97$, which gives observations 2 years ago half the weight as current observations for purposes of calculating $\bar{m}_{X,t}$. The resulting values of $\bar{m}_{X,t}$ are plotted in the first three panels of Figure 11. Both \bar{m}_{Nt} and \bar{m}_{Et} rise over time, while $\bar{m}_{U,t}$ is countercyclical without exhibiting a particular trend. The secular rise in \bar{m}_{Nt} and \bar{m}_{Et} suggests that the upward trend in missing individuals likely comes from N and E . The countercyclical

³³Our approach thus allocates an average fraction $m = m_E + m_N + m_U = 0.1567$ of the M to a status E , N , or U ; the vast majority of M are not allocated to any status. A fraction $m_E/m = 0.6222$ of those allocated are designated as E and fractions $m_N/m = 0.3031$ and $m_U/m = 0.0744$ designated as N and U , respectively. These compare with fractions $\pi_E^*/(1 - \pi_M^*) = 0.6141$ of individuals who are originally either E , N or U who were reported to be employed and fractions $\pi_N^*/(1 - \pi_M^*) = 0.3427$ and $\pi_U^*/(1 - \pi_M^*) = 0.0433$ who were N or U . Thus our adjustment for missing observations raises the count of U , lowers the count of N , and does not much change the count of E . The reason is that MU transitions are more common and MN transitions less common than they would be if the population of M the previous month had the same characteristics as those for whom a status E , N or U was observed.

behavior of \bar{m}_{Ut} tells us that unemployed individuals are more likely to be missed during a weak labor market.

Other panels of Figure 11 plot month-by-month estimates of some of the parameters whose full-sample maximum likelihood estimates were reported in Table 3. These were found by fixing the digit-preference parameters θ_A at their full-sample averages and then maximizing the likelihood of observation t alone with respect to the other parameters in Table 3. These in turn were exponentially smoothed. The fractions of NU and MU transitions that individuals perceive as continuations of long-term unemployment ($\bar{q}_{6,NU,t}$ and $\bar{q}_{6,MU,t}$) rose sharply during the Great Recession and have been slow to return to their historical averages. Both the perceived weekly UU continuation probability for type 1 individuals \bar{p}_{1t} and the objective monthly probability $\bar{\gamma}_{1,UU,t}$ react to seasonal hiring, consistent with the high seasonality in unadjusted short-term unemployment, and both fell during the Great Recession.³⁴ For type 2 individuals, there is a time trend in perceived \bar{p}_{2t} that is not fully matched by that for the objective $\bar{\gamma}_{2,UU,t}$ probability, though both increased significantly in the Great Recession and were slow to come down afterward. The fraction \bar{w}_{2t} of type 2 workers among the reported unemployed rose through 2011 and has been slowly declining since.

4.5 Reconciled estimates of labor-force participation and unemployment rates.

Our reconciled estimates for labor-force status in month $t - 1$ begin with the value of π_{t-1}^* calculated as described in Section 4.1, which estimates labor-force status as it would be measured using the first-interview technology. We then adjust this using $\bar{m}_{X,t-1}$ described in Section 4.3 based on individuals who were missing in $t - 1$ but for whom one of the statuses E , N , or U was reported in t . Next we adjust N_{t-1} down and U_{t-1} up based on the number of individuals who reported status N in month $t - 1$ and reported in month t that they were unemployed and had been looking for work for longer than 4 weeks. The fraction of individuals for whom this adjustment is

³⁴The feature of the data that gives rise to this conclusion is the observation that individuals with unemployment durations of 5-14 weeks were much more likely to remain unemployed during the Great Recession, meaning that more type 1 individuals must have exited unemployment after just one month of unemployment. One possible interpretation is that individuals would only voluntarily quit their job in this episode if they knew they could get another job quickly. A drop in p_{1t} during the Great Recession was also found by Ahn and Hamilton (2019, Figure 4 and Table 1). They found that this feature was unique to the Great Recession and was not seen in other recessions.

warranted in month $t - 1$ is given by

$$m_{N,t-1}^{\sharp} = \frac{\sum_{\tau=5}^{99} \sum_{j \in J} y_{N,U,t}^{[j]}(\tau)}{\sum_{j \in J} \left[y_{E,t-1}^{[j-1]} + y_{N,t-1}^{[j-1]} + y_{M,t-1}^{[j-1]} + y_{U,t-1}^{[j-1]} \right]}.$$

This averages 0.37% of all individuals over the full sample, so it is quite a significant adjustment. We then further adjust N_{t-1} and U_{t-1} based on reinterpreting a fraction of the $U_{t-2}^{15+} N_{t-1}$ transitions as UU continuations. We construct monthly estimates of $m_{N,t-1}^b$, the fraction of the population with reported UN who are better interpreted as long-term UU , from

$$m_{N,t-1}^b = \pi_{U,t-1}^* \bar{w}_{2,t-1} \bar{\gamma}_{2,UN,t-1} \xi_{UN}.$$

Here $\pi_{U,t-1}^*$ is the fourth element of π_{t-1}^* , $\bar{w}_{2,t-1}$ and $\bar{\gamma}_{2,UN,t-1}$ are the exponentially smoothed parameters plotted in panels 8 and 10 of Figure 11, and we fix $\xi_{UN} = 0.627$ at the full-sample average³⁵.

The adjustments $m_{N,t-1}^{\sharp}$ and $m_{N,t-1}^b$ entail some double-counting of individuals who are $U_{t-2}^{15+} N_{t-1} U_t^{5+}$ who would be included in both $m_{N,t-1}^{\sharp}$ and $m_{N,t-1}^b$. We correct for this by calculating k^{\natural} , the fraction of $m_{N,t}^{\sharp} + m_{N,t}^b$ that comes from double-counting the same individuals, from our full-sample estimate of that fraction:

$$k^{\natural} = \frac{m_N^{\natural}}{m_N^{\sharp} + m_N^b} = \frac{0.0006}{0.0037 + 0.0026} = 0.095$$

giving rise to the monthly estimate $m_{N,t}^{\natural} = k^{\natural}(m_{N,t}^{\sharp} + m_{N,t}^b)$. Our final estimates that correct for rotation bias, non-randomly missing observations, and misclassified N are then

$$\begin{bmatrix} \tilde{\pi}_{E,t-1} \\ \tilde{\pi}_{N,t-1} \\ \tilde{\pi}_{M,t-1} \\ \tilde{\pi}_{U,t-1} \end{bmatrix} = \begin{bmatrix} \pi_{E,t-1}^* + \pi_{M,t-1}^* \bar{m}_{E,t-1} \\ \pi_{N,t-1}^* + \pi_{M,t-1}^* \bar{m}_{N,t-1} - m_{N,t-1}^{\sharp} - m_{N,t-1}^b + m_{N,t-1}^{\natural} \\ \pi_{M,t-1}^* (1 - \bar{m}_{E,t-1} - \bar{m}_{N,t-1} - \bar{m}_{U,t-1}) \\ \pi_{U,t-1}^* + \pi_{M,t-1}^* \bar{m}_{U,t-1} + m_{N,t-1}^{\sharp} + m_{N,t-1}^b - m_{N,t-1}^{\natural} \end{bmatrix}. \quad (35)$$

³⁵We obtained similar results allowing $\xi_{UN,t}$ to change over time.

Our adjusted estimates of the unemployment rate and labor-force participation rate are

$$\tilde{u}_t = \tilde{\pi}_{U,t} / (\tilde{\pi}_{E,t} + \tilde{\pi}_{U,t})$$

$$\tilde{\ell}_t = (\tilde{\pi}_{E,t} + \tilde{\pi}_{U,t}) / (\tilde{\pi}_{E,t} + \tilde{\pi}_{N,t} + \tilde{\pi}_{U,t}).$$

Note that these are all seasonally unadjusted magnitudes in order to preserve all the accounting identities associated with observed transitions. To relate these to the usually reported magnitudes, we plotted seasonally-adjusted values for these rates in Figures 1 and 12.³⁶

The black lines in Figure 12 show the BLS values for the unemployment rate and labor-force participation rate, and the first row of Table 10 reports their values over the full sample. We can calculate the effect of our correction for rotation bias alone by setting $\bar{m}_{Et} = \bar{m}_{Nt} = \bar{m}_{Ut} = m_{Nt}^{\#} = m_{Nt}^b = 0$ in (35). These series are plotted as the red lines in Figure 12, with the full-sample average reported in the second row of Table 10. Correcting for rotation bias alone would add half a percentage point to the unemployment rate and 1.1% to the labor-force participation rate. The green lines in Figure 12 and third row of Table 10 show the contribution of also taking account of the nonrandom nature of missing observations (that is, allows for nonzero $\bar{m}_{Et}, \bar{m}_{Nt}, \bar{m}_{Ut}$). The blue lines and last row of Table 10 show the effects of all three adjustments. Altogether, our adjustments add 2.0% to the unemployment rate and 2.1% to the labor-force participation rate on average. For the unemployment rate, the NU misclassification is the main source of cyclical features in the errors. For the labor-force-participation rate, both rotation bias and missing observations explain the slowly rising trend in the errors and the NU misclassification explains the countercyclicality.

The last column of Table 10 shows that while rotation bias matters for the employment-population ratio, the ratio is unchanged after correcting for missing observations or misclassified N . Thus the employment-population ratio could be a more robust measure of the labor-market slack in the presence of increasing nonresponses and errors in responses in the CPS.

The top panel of Figure 13 compares our adjusted estimate \tilde{u}_t (in blue) with three different unemployment rates reported by the BLS— the usual U3 unemployment rate (black) along with U5 unemployment (red), which includes discouraged workers and all other marginally attached workers, and U6 unemployment (green) which adds people who are employed part-time for economic reasons.

³⁶These were calculated using the X11 instruction in RATS.

Our adjustment includes more individuals than U5, but far less than U6.

4.6 Reconciled estimates of unemployment-continuation probabilities.

Our concept for calculating unemployment-continuation probabilities is that used by Fujita and Ramey (2009) and Elsby, Hobijn and Şahin (2010) – we track the objective labor-force status next month of someone who is unemployed this month. However, our estimates differ from theirs in that we correct for rotation bias, nonrandom missing observations, and misclassification of some N .

Let $\bar{\gamma}_{i,U,t}$ be the (4×1) vector of smoothed transition probabilities for unemployed individuals of type i in month t . An individual element of the vector $\bar{\gamma}_{i,UX,t}$ represents the probability that a type i individual who was reported to be unemployed in rotation 1 or 5 would have reported status $X \in \{E, N, M, U\}$ in rotation 2 or 6. Adjusting this to correct for rotation bias is achieved by

$$\bar{\gamma}_{i,U,t}^* = (1/2)(\bar{R}_t^{[2]} + \bar{R}_t^{[6]})\bar{\gamma}_{i,U,t}.$$

We further concluded that a fraction ξ_{UN} of the type 2 individuals who report N in $t + 1$ should be viewed as UU continuations. Correcting for missing observations, this gives an estimate of the true unemployment-continuation probability for type 2 individuals of

$$\hat{\gamma}_{2,UU,t}^* = \frac{\bar{\gamma}_{2,UU,t}^* + \xi_{UN}\bar{\gamma}_{2,UN,t}^*}{1 - \bar{\gamma}_{2,UM,t}^*}.$$

This series is plotted as the red line in the last panel of Figure 11. We calculate monthly unemployment-continuation probabilities for type 1 individuals from $\hat{\gamma}_{1,UU,t}^* = \bar{\gamma}_{1,UU,t}^*/(1 - \bar{\gamma}_{1,UM,t}^*)$.

The estimate $\hat{\gamma}_{2,UU}^*$ averages 0.77, well below $p_2^{4.33} = 0.89$, the value we would have expected based on reported unemployment durations. Nevertheless, the adjustment goes a fair way toward reconciling perceived durations with objective continuation probabilities. One source of the remaining discrepancy between our estimate of the objective continuation probability $\hat{\gamma}_{2,UU}^*$ and the perceived duration of job search p_2 is on-the-job search. Recall from Section 3.2 that EU^{5+} transitions account for 26% of EU observations, with many EU individuals reporting duration longer than 6 months. As noted by Kudlyak and Lange (2018), we could interpret these individuals as correctly reporting how long they have been looking for a job or looking for a better job, while still

defending the estimate $\tilde{\gamma}_{2,UU}^*$ as a correct summary of the true probability of remaining unemployed without an intervening spell of employment. A second possible source of discrepancy between $\tilde{\gamma}_{2,UU}^*$ and p_2 is that individuals are reporting not the length of a continuous spell of unemployment but instead how long it has been since their last good job (Elsby et al. (2011); Farber and Valletta (2015)). We conclude that our procedure of adjusting unemployment-continuation probabilities up, but not all the way to those implied by reported job-search durations, is the correct way to reconcile the data.

We next calculate the fraction $\tilde{w}_{i,t-1}$ of total unemployed individuals $\tilde{\pi}_{U,t-1}$ that are of type i . Consider the last row of equation (35). For the first term in that equation ($\pi_{U,t-1}^*$), we know the fraction of type i from the estimate of $\bar{w}_{i,t-1}$. We assume the same fraction $\bar{w}_{i,t-1}$ could be used to impute types to the missing unemployed for the second term ($\pi_{M,t-1}^* \bar{m}_{U,t-1}$). The third term ($m_{N,t-1}^\#$) is derived from observed NU^{5+} transitions, for which we have estimated the fraction of type 1 to be $\bar{q}_{5,NU,t-1}/(\bar{q}_{5,NU,t-1} + \bar{q}_{6,NU,t-1})$. The last two terms by construction come solely from type 2 individuals. We thus estimate

$$\tilde{w}_{1,t} = \frac{\bar{w}_{1,t}(\pi_{U,t}^* + \pi_{M,t}^* \bar{m}_{U,t}) + m_{N,t}^\# \bar{q}_{5,NU,t}/(\bar{q}_{5,NU,t} + \bar{q}_{6,NU,t})}{\tilde{\pi}_{U,t}}$$

and $\tilde{w}_{2,t} = 1 - \tilde{w}_{1,t}$. Our estimate of the true monthly continuation probability averaged across all individuals who are truly unemployed is then

$$\tilde{w}_{1,t} \tilde{\gamma}_{1,UU,t}^* + \tilde{w}_{2,t} \tilde{\gamma}_{2,UU,t}^*,$$

which is the series plotted as the blue line in Panel A of Figure 1.

4.7 Reconciled estimates of new flows into unemployment.

We estimate that a fraction $\tilde{w}_{i,t} \tilde{\pi}_{U,t}$ of individuals in the sample are truly unemployed of type $i \in \{1, 2\}$ in month t . Of these, a fraction $\tilde{\gamma}_{i,UU,t+1}$ are still unemployed the next month, giving rise to

$$\tilde{V}_{i,t+1} = \tilde{w}_{i,t+1} \tilde{\pi}_{U,t+1} - \tilde{\gamma}_{i,UU,t+1} \tilde{w}_{it} \tilde{\pi}_{U,t} \quad (36)$$

as an estimate of the number of individuals of type i who are newly unemployed in month $t + 1$ and $\tilde{V}_{t+1} = \tilde{V}_{1,t+1} + \tilde{V}_{2,t+1}$ as the total number of newly unemployed. This is the series that was plotted as the blue line in Panel B of Figure 1. We also reproduce it as the blue line in the bottom panel of Figure 13 along with several alternative estimates. Shimer (2012) and other researchers have estimated unemployment inflows from the number of unemployed with reported durations of less than 5 weeks, shown in red as a percent of the civilian population. Others like Fujita and Ramey (2012) base their calculation on the number of EU and NU transitions among those with two consecutive months of nonmissing observations,

$$\hat{V}_t = \frac{\sum_{j \in J} \left(y_{E,U,t}^{[j]} + y_{N,U,t}^{[j]} \right)}{\sum_{j \in J} \left(y_{E,E,t}^{[j]} + y_{E,N,t}^{[j]} + y_{E,U,t}^{[j]} + y_{N,E,t}^{[j]} + y_{N,N,t}^{[j]} + y_{N,U,t}^{[j]} + y_{U,E,t}^{[j]} + y_{U,N,t}^{[j]} + y_{U,U,t}^{[j]} \right)}, \quad (37)$$

shown as the turquoise line. The Shimer estimate is significantly below the Fujita-Ramey estimate because the latter includes EU^{5+} and NU^{5+} transitions. Our estimate is above \hat{V}_t . The biggest single reason for this is rotation bias, which causes flows into unemployment as calculated from the numerator of (37) to be smaller than flows out of unemployment even in months when the measured unemployment rate is constant or even rising. One can see the effect of rotation bias by replacing $\sum_{j \in J} y_{X_1,X_2,t}^{[j]}$ in (37) by the estimate $\pi_{X_1,t-1}^* \pi_{X_1,X_2,t}^*$. This corrects the calculation for rotation bias but makes no other adjustments. The resulting series \hat{V}_t^* is shown as the green line in Figure 13, which is much higher than the estimate \hat{V}_t from (37). Our fully adjusted series \tilde{V}_t makes a number of other adjustments that can either increase or decrease the estimate relative to \hat{V}_t^* . We exclude NU^{5+} transitions because we see them as continuing spells of unemployment, which lowers the estimate of V . But we also adjust the estimate up as a result of our treatment of missing observations. On average \tilde{V}_t is above \hat{V}_t^* , but rotation bias is the biggest single problem with \hat{V}_t . Finally, we note that the BLS also publishes estimates of the number of EU and NU flows that are consistent with observed stocks of E , N and U . Their series (shown in black) adjusts the data in the direction of our estimates (that is, it is above \hat{V}_t) but is lower than an adjustment that only corrects for rotation bias (the BLS estimate is below \hat{V}_t^*).

4.8 Reconciled estimates of unemployment duration.

Let $\tilde{V}_{i,t-d+1}$ denote the number of newly unemployed of type i at $t - d + 1$ as calculated in (36). A fraction $\tilde{\gamma}_{i,UU,t-d+2}$ will still be unemployed at $t - d + 2$. Thus the number unemployed for exactly d months as of month t would be given by

$$\tilde{U}_{i,t}^d = \tilde{V}_{i,t-d+1} \tilde{\gamma}_{i,UU,t-d+2} \cdots \tilde{\gamma}_{i,UU,t-2} \tilde{\gamma}_{i,UU,t-1} \tilde{\gamma}_{i,UU,t}. \quad (38)$$

This implies an average unemployment duration of those who are unemployed in month t of

$$\tilde{d}_t = \frac{\sum_{d=1}^{48} d(\tilde{U}_{1,t}^d + \tilde{U}_{2,t}^d)}{\sum_{d=1}^{48} (\tilde{U}_{1,t}^d + \tilde{U}_{2,t}^d)}. \quad (39)$$

Dividing by 4.33 gives the unemployment duration in weeks plotted as the blue lines in Panel E of Figure 1. Our series is much lower on average and less cyclically variable than the BLS measure in black.

We calculated the average values of (38) over all months t in our sample and report in Table 11 the average percentage of the truly unemployed $\tilde{\pi}_U$ for whom the true duration is less than 5 weeks (1 month), 5-14 weeks (2-3 months), 15-26 weeks (4-6 months) and longer than 26 months (7 months and over), along with the average duration. Our estimate of the average duration of unemployment is only 15 weeks, about 11 weeks lower than the BLS reports. Kudlyak and Lange (2018) constructed estimates of the number of newly unemployed as a fraction of total unemployed by (1) counting all $E_{t-1}U_t$ as newly unemployed despite the duration of search reported at t , and (2) also counting all $N_{t-1}U_t$ as newly unemployed. Our estimate of the fraction of individuals unemployed for less than 5 weeks, 36.4%, is in between their two estimates (29.1% and 46.1%, respectively) because we designate some, but not all, of the $N_{t-1}U_t$ as unemployed at $t - 1$. Their two methods produced estimates of 37.5% and 24.1%, respectively, for the fraction of unemployed with duration greater than 14 weeks, with our estimate of 32% again in between those two. Although their approach did not allow them to uncover the average duration of unemployment, their calculations confirm our conclusion that the BLS estimates substantially overstate the number of long-term unemployed.

5 Conclusion.

The data underlying the CPS contain multiple internal inconsistencies. These include the facts that people's answers change the more times they are asked the same question, stock estimates are inconsistent with flow estimates, missing observations are not random, reported unemployment durations are inconsistent with reported labor-force histories, and people prefer to report some numbers over others. Ours is the first paper to attempt a unified reconciliation of these issues. We conclude that the U.S. unemployment rate and labor-force continuation rates are higher than conventionally reported while the average duration of unemployment is considerably lower.

References

- Abowd, John M., and Arnold Zellner (1985). "Estimating Gross Labor-Force Flows." *Journal of Business and Economic Statistics* 3, no. 3: 254-283.
- Abraham, Katharine G., and Shimer Robert (2001). "Changes in Unemployment Duration and Labor Force Attachment." NBER Working Paper 8513.
- Ahn, Hie Joo, and James D. Hamilton (2020). "Heterogeneity and Unemployment Dynamics." *Journal of Business and Economic Statistics*, forthcoming.
- Ahn, Hie Joo, and Ling Shao (2017). "Precautionary On-the-Job Search over the Business Cycle." Federal Reserve Finance and Economics Discussion Series 2017-025.
- Bailar, Barbara A (1975). "The Effects of Rotation Group Bias on Estimates from Panel Surveys." *Journal of the American Statistical Association* 70: 23-30.
- Baker, Michael (1992). "Digit Preference in CPS Unemployment Data." *Economics Letters*, 39(1):117-121.
- Biemer, Paul P., and John M. Bushery (2000). "On the Validity of Markov Latent Class Analysis for Estimating Classification Error in Labor Force Data." *Survey Methodology* 26: 139-152.
- Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin (2010). "The Labor Market in the Great Recession." *Brookings Papers on Economic Activity*, Spring 2010: 1-56.
- Elsby, Michael W.L., Bart Hobijn, and Ayşegül Şahin (2015). "On the Importance of the Participation Margin for Labor Market Fluctuations." *Journal of Monetary Economics* 72: 64-82.
- Elsby, Michael WL, Bart Hobijn, Ayşegül Şahin, and Robert G. Valletta (2011). "The Labor Market in the Great Recession— An Update to September 2011." *Brookings Papers on Economic*

Activity Fall 2011: 353-371.

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

Faberman, R. Jason, Andreas I. Mueller, Aysegül Şahin, and Giorgio Topa (2019). "Job Search Behavior among the Employed and Non-Employed." working paper, University of Texas at Austin.

Farber, Henry S., and Robert G. Valletta (2015). "Do Extended Unemployment Benefits Lengthen Unemployment Spells? Evidence from Recent Cycles in the US Labor Market." *Journal of Human Resources* 50: 873-909.

Feng, Shuaizhang, and Yingyao Hu (2013). "Misclassification Errors and the Underestimation of the US Unemployment rate." *American Economic Review* 103: 1054-70.

Fujita, Shigeru and Garey Ramey (2009). "The Cyclicalities of Separation and Job Finding Rates." *International Economic Review*, 50(2): 415-430.

Hall, Robert E., and Marianna Kudlyak (2019). "Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics." NBER Working Paper 25625.

Halpern-Manners, Andrew, and John Robert Warren (2012). "Panel Conditioning in Longitudinal Studies: Evidence from Labor Force Items in the Current Population Survey." *Demography* 49(4): 1499-1519.

Hamilton, James D. (1994). *Time Series Analysis*. Princeton: Princeton University Press.

Hansen Morris H., William N. Hurwitz, Harold Nisselson, and Joseph Steinberg (1955). "The Redesign of the Census Current Population Survey." *Journal of the American Statistical Association* 50: 701-719.

Hirsch, Barry T., and John V. Winters (2016). "Rotation Group Bias in Measures of Multiple Job Holding." *Economics Letters* 147: 160-163.

Krueger, Alan B., Alexandre Mas, and Xiaotong Niu (2017). "The Evolution of Rotation Group Bias: Will the Real Unemployment Rate Please Stand Up?" *Review of Economics and Statistics* 99: 258-264.

Kudlyak, Marianna, and Fabian Lange (2018). "Measuring Heterogeneity in Job Finding Rates Among the Nonemployed Using Labor Force Status Histories." Working paper, Federal Reserve Bank of San Francisco.

Madrian, Brigitte C., and Lars John Lefgren (2000). "An Approach to Longitudinally Matching

Current Population Survey (CPS) Respondents." *Journal of Economic and Social Measurement* 26(1): 31-62.

Meyer, Bruce D., Wallace K. C. Mok, and James X. Sullivan (2015). "Household Surveys in Crisis." *Journal of Economic Perspectives*, 29(4): 199-226.

Nekarda, Christopher J. (2009). "A Longitudinal Analysis of the Current Population Survey: Assessing the Cyclical Bias of Geographic Mobility." Federal Reserve Board of Governors.

Rothstein, Jesse (2011). "Unemployment Insurance and Job Search in the Great Recession." *Brookings Papers on Economic Activity*, Fall 2011: 143-196.

Ryu, Hang K., and Daniel J. Slottje (2000). "Estimating the Density of Unemployment Duration Based on Contaminated Samples or Small Samples." *Journal of Econometrics*, 95(1): 131-156.

Shibata, Ippei (2019). "Reassessing Classification Errors in the Analysis of Labor Market Dynamics." IMF Working Paper.

Shimer, Robert (2012). "Reassessing the Ins and Outs of Unemployment." *Review of Economic Dynamics*, 15(2): 127-148.

Silverstone, Brian, and Will Bel (2010). "Labour Market Flows in New Zealand: Some Questions and Some Answers." 51st Conference of the New Zealand Association of Economists, Auckland, vol. 30.

Solon, Gary (1986). "Effects of Rotation Group Bias on Estimation of Unemployment." *Journal of Business and Economic Statistics* 4: 105-109.

Torelli, Nicola, and Ugo Trivellato (1993). "Modelling Inaccuracies in Job-search Duration Data." *Journal of Econometrics* 59(1-2): 187-211.

Van den Berg, Gerald J., and Bas van der Klaauw (2001). "Combining Micro and Macro Unemployment Duration Data." *Journal of Econometrics* 102: 271-309.

Van den Brakel, Jan A., and Sabine Krieg (2015). "Dealing with Small Sample Sizes, Rotation Group Bias and Discontinuities in a Rotating Panel Design." *Survey Methodology* 41: 267-296.

Table 1. Parameters estimated separately for rotation 1, rotation 5, and NX, EX and MX transitions from rotation 1 to rotation 2.

	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]
param	rotation 1 only	std error	rotation 5 only	std error	NX only	std error	EX only	std error	MX only	std error
p_1	0.8261	0.0031	0.8272	0.0022	0.7529	0.0101	0.7534	0.0115	0.8303	0.0046
p_2	0.9737	0.0025	0.9733	0.0025	0.9745	0.0021	0.9688	0.0033	0.9741	0.0024
w_1	0.4320	0.0471	0.4081	0.0481						
π_E	0.4253		0.4203							
π_N	0.2365	0.0051	0.2481	0.0052						
π_M	0.3083	0.0038	0.3047	0.0039						
π_U	0.0300	0.0031	0.0269	0.0029						
π_{XE}					0.0384	0.0015	0.8881	0.0043	0.1300	0.0053
π_{XN}					0.8762	0.0007	0.0319	0.0006	0.0664	0.0036
π_{XM}					0.0609		0.0669		0.7933	
π_{XU}					0.0246	0.0020	0.0130	0.0008	0.0103	0.0008
q_1					0.0935		0.2185		0.0906	
q_2					0.0776	0.0055	0.1896	0.0059	0.0995	0.0072
q_3					0.0825	0.0060	0.1798	0.0040	0.0834	0.0071
q_4					0.0534	0.0028	0.1214	0.0049	0.0800	0.0043
q_5					0.1880	0.0198	0.1197	0.0029	0.2237	0.0222
q_6					0.5049	0.0411	0.1709	0.0149	0.4229	0.0482
$q_5 + q_6$					0.6929		0.2906		0.6465	
$\theta_{A,1}$	0.1254	0.0036	0.1295	0.0077	0.2095	0.0285	0.1110	0.0262	0.0769	0.0260
$\theta_{A,2}$	0.7740	0.0029	0.7407	0.0029	0.7499	0.0066	0.7048	0.0194	0.7328	0.0108
$\theta_{A,3}$	0.4840	0.0089	0.4616	0.0105	0.4917	0.0173	0.3900	0.0306	0.5168	0.0228
$\theta_{A,4}$	0.9255	0.0035	0.8623	0.0111	0.8597	0.0114	0.8366	0.0304	0.8433	0.0227
$\theta_{A,5}$	0.7248	0.0145	0.6893	0.0091	0.7062	0.0156	0.6907	0.0308	0.7775	0.0423
$\theta_{A,6}$	0.9229	0.0088	0.9088	0.0045	0.8749	0.0146	0.8243	0.0167	0.8890	0.0191
$\theta_{A,7}$	0.9601	0.0080	0.9592	0.0022	0.9508	0.0142	0.9716	0.0102	0.9362	0.0174
$\theta_{A,8}$	0.8956	0.0067	0.8619	0.0101	0.7310	0.0356	0.7445	0.0245	0.7957	0.0332
$\theta_{A,9}$	0.9410	0.0104	0.9492	0.0039	0.9502	0.0125	0.9500	0.0342	0.9347	0.0101
$\theta_{A,10}$	0.1654	0.0083	0.1511	0.0062	0.1990	0.0099	0.1367	0.0099	0.1477	0.0080
$\theta_{A,11}$	0.4930	0.0081	0.4983	0.0043	0.5892	0.0097	0.4986	0.0232	0.4871	0.0157
$\theta_{A,12}$	0.8978	0.0164	0.8947	0.0129	0.9222	0.0093	0.9106	0.0185	0.9084	0.0082
$\theta_{A,13}$	0.1643	0.0279	0.1096	0.0397	0.1257	0.0238	0.1442	0.0123	0.0879	0.0445

Table 2. Parameters estimated separately for UX transitions from rotations 1 to 2 and 5 to 6.

		$\gamma_{1,UE}$	$\gamma_{1,UN}$	$\gamma_{1,UM}$	$\gamma_{1,UU}$	$\gamma_{2,UE}$	$\gamma_{2,UN}$	$\gamma_{2,UM}$	$\gamma_{2,UU}$
[1]	Rotation 1 estimate	0.3189	0.2176	0.0928	0.3707	0.1164	0.2376	0.0743	0.5718
[2]	Standard error	0.0057	0.0027	0.0037	...	0.0093	0.0099	0.0033	...
[3]	Rotation 5 estimate	0.3391	0.2181	0.0898	0.3530	0.1214	0.2240	0.0695	0.5851
[4]	Standard error	0.0080	0.0018	0.0017	...	0.0079	0.0075	0.0041	...

Table 3. Parameters estimated jointly across all rotations.

	estimate		estimate		estimate		estimate		estimate		estimate
p_1	0.8092	$\theta_{A,6}$	0.8890	$q_{1,EU}$	0.2250	$q_{1,NU}$	0.0876	$q_{1,MU}$	0.1011	$\gamma_{1,UE}$	0.3283
p_2	0.9733	$\theta_{A,7}$	0.9547	$q_{2,EU}$	0.1873	$q_{2,NU}$	0.0820	$q_{2,MU}$	0.0967	$\gamma_{1,UN}$	0.2179
w_1	0.4003	$\theta_{A,8}$	0.8030	$q_{3,EU}$	0.1977	$q_{3,NU}$	0.0771	$q_{3,MU}$	0.0984	$\gamma_{1,UM}$	0.0915
$\theta_{A,1}$	0.1453	$\theta_{A,9}$	0.9394	$q_{4,EU}$	0.1141	$q_{4,NU}$	0.0655	$q_{4,MU}$	0.0702	$\gamma_{1,UU}$	0.3624
$\theta_{A,2}$	0.7357	$\theta_{A,10}$	0.1711	$q_{5,EU}$	0.1369	$q_{5,NU}$	0.1861	$q_{5,MU}$	0.2141	$\gamma_{2,UE}$	0.1188
$\theta_{A,3}$	0.4693	$\theta_{A,11}$	0.5215	$q_{6,EU}$	0.1390	$q_{6,NU}$	0.5017	$q_{6,MU}$	0.4196	$\gamma_{2,UN}$	0.2311
$\theta_{A,4}$	0.8743	$\theta_{A,12}$	0.9034							$\gamma_{2,UM}$	0.0720
$\theta_{A,5}$	0.7109	$\theta_{A,13}$	0.1194							$\gamma_{2,UU}$	0.5782

Notes to Table 3. Also estimated (but not reported) are separate coefficients $\pi_{XE}, \pi_{XN}, \pi_{XM}, \pi_{XU}$ for $X \in \{E, N, M\}$.

Table 4. Average numbers of individuals with indicated status across different rotation groups.

	[1]	[2]	[3]	[4]	[5]	[6]	[7]
rotation	E	N	M	U	total	U/(U+E)	(U+E)/(U+E+N)
1	7,837	4,357	5,621	553	18,368	6.6	65.8
2	7,982	4,573	5,307	539	18,401	6.3	65.1
3	7,988	4,619	5,295	520	18,422	6.1	64.8
4	7,971	4,636	5,324	507	18,438	6.0	64.6
5	7,773	4,588	5,621	497	18,479	6.0	64.3
6	7,883	4,675	5,423	489	18,470	5.8	64.2
7	7,914	4,692	5,373	480	18,459	5.7	64.1
8	7,963	4,719	5,277	482	18,441	5.7	64.2

Table 5. Values of rotation-group bias parameters for full sample.

j	1	2	3	4	5	6	7	8
$\theta_{EM}^{[j]}$	0	0.0176	0.0180	0.0155	0	0.0021	0.0055	0.0113
$\theta_{NU}^{[j]}$	0	0.0031	0.0072	0.0099	0.0126	0.0140	0.0159	0.0154
$\theta_{NM}^{[j]}$	0	0.0437	0.0486	0.0490	0.0343	0.0506	0.0515	0.0569

Table 6. Estimated average fractions of individuals π_X^* who would have reported labor status E, N, M, U and transition probabilities π_{X_1, X_2}^* if all individuals were being interviewed for the first time.

$$\begin{bmatrix} \pi_E^* \\ \pi_N^* \\ \pi_M^* \\ \pi_U^* \end{bmatrix} = \begin{bmatrix} 0.4240 \\ 0.2366 \\ 0.3096 \\ 0.0299 \end{bmatrix} \quad \begin{bmatrix} \pi_{EE}^* & \pi_{NE}^* & \pi_{ME}^* & \pi_{UE}^* \\ \pi_{EN}^* & \pi_{NN}^* & \pi_{MN}^* & \pi_{UN}^* \\ \pi_{EM}^* & \pi_{NM}^* & \pi_{MM}^* & \pi_{UM}^* \\ \pi_{EU}^* & \pi_{NU}^* & \pi_{MU}^* & \pi_{UU}^* \end{bmatrix} = \begin{bmatrix} 0.8984 & 0.0364 & 0.0917 & 0.2033 \\ 0.0257 & 0.8684 & 0.0460 & 0.2003 \\ 0.0636 & 0.0662 & 0.8538 & 0.0880 \\ 0.0123 & 0.0290 & 0.0085 & 0.5084 \end{bmatrix}$$

Table 7. Characteristics of U and N as a function of rotation.

	1	2	3	4	5	6	7	8	avg(2-8)
(1) $N/(E+N+U)$	34.2	34.9	35.2	35.4	35.7	35.8	35.9	35.8	35.5
(2) retired/ $(E+N+U)$	15.5	16.2	16.4	16.6	16.2	16.6	16.8	16.8	16.5
(3) disabled/ $(E+N+U)$	4.6	5.0	5.2	5.3	4.9	5.2	5.3	5.4	5.2
(4) Probability E or U in $j+1$ given retired or disabled in j	1.75	1.90	1.87		1.91	1.82	1.81		1.86
(5) Standard error	(0.017)	(0.017)	(0.017)		(0.018)	(0.017)	(0.017)		(0.006)
(6) U (self-report)/ $(E+U)$	3.3	3.1	3.0	2.9	3.0	2.9	2.8	2.8	2.9
(7) U (proxy)/ $(E+U)$	3.2	3.1	3.0	3.0	3.0	2.9	2.9	2.9	3.0
(8) $E/(E+U+N)$	61.5	60.9	60.8	60.8	60.4	60.3	60.4	60.4	60.6
(9) $\dot{E}/(E+U+N)$	16.6	16.1	15.9	15.8	15.6	15.5	15.5	15.5	15.7
(10) $U/(E+U)$	6.6	6.4	6.2	6.0	6.0	5.9	5.8	5.7	
(11) $U/(E+U)$ given M1		7.4	6.8	6.3	5.6	5.4	5.2	5.1	
(12) $U/(E+U)$ given M1 and M2			8.8	8.1	6.5	6.3	5.9	5.9	
(13) $U/(E+U)$ given M1-M3				9.9	6.9	6.7	6.3	6.5	
(14) $U/(E+U)$ given M1-M4					9.3	8.6	7.9	7.6	
(15) $U/(U+E)$ given M1-M5						9.2	8.5	7.8	
(16) $U/(U+E)$ given M1-M6							10.2	9.0	
(17) Total N (in thousands)	976	1024	1035	1038	1028	1047	1051	1057	

Notes to Table 7. Row (1): N as a percent of $E+N+U$. (2): retired individuals as a percent of $E+N+U$. (3): disabled individuals as a percent of $E+N+U$. (4): probability that an individual who is retired or disabled in rotation j will be E or U in rotation $j+1$. (5): standard error of row (4). (6): individuals who report their own status to be U as a percent of the labor force. (7): individuals whose status is reported by another member of the household to be U as a percent of the labor force. (8): E as a percent of $E+N+U$. (9): Part-time unemployed plus those usually employed but not employed this week as a percent of $E+N+U$. (10): unemployment rate as a function of rotation among individuals who are not missing in rotation 1. (11): unemployment rate as a function of rotation among individuals who are missing in rotation 1 but not missing in rotation 2. (12): unemployment rate among individuals who are missing in rotations 1 and 2 but not missing in 3. Rows (13)-(16): unemployment rate among individuals who are missing in rotations 1 through $j-1$ but not missing in j . (17): Total number of individuals counted as not in the labor force from each rotation. All numbers are reported as percent except for last row which is in thousands of individuals. Rows (1)-(9) and (17) refer to average over Jul 2001 to Feb 2020 while rows (8)-(14) are over Sep 2002 to Feb 2020.

Table 8. Unemployment rates in rotation j and $j + 1$ among individuals who are not missing in either j or $j + 1$.

Rotation	u_j	u_{j+1}	difference
$j = 1$	6.50	6.23	0.27
$j = 2$	6.22	5.99	0.24
$j = 3$	6.00	5.85	0.15
$j = 5$	5.91	5.71	0.19
$j = 6$	5.73	5.58	0.15
$j = 7$	5.62	5.59	0.04

Table 9. Month $t + 1$ employment probabilities for UUU and UNU histories.

UUU	Probability	UNU	Probability
$U_{t-2}^{1.4}, U_{t-1}^{5.14}, U_t^{5.14}$	0.19	$U_{t-2}^{1.4}, N_{t-1}, U_t^{5.14}$	0.15
$U_{t-2}^{5.14}, U_{t-1}^{5.14}, U_t^{15.26}$	0.16	$U_{t-2}^{5.14}, N_{t-1}, U_t^{15.26}$	0.15
$U_{t-2}^{15.26}, U_{t-1}^{15.26}, U_t^{15.26}$	0.14	$U_{t-2}^{15.26}, N_{t-1}, U_t^{15.26}$	0.14
$U_{t-2}^{15.26}, U_{t-1}^{27+}, U_t^{27+}$	0.12	$U_{t-2}^{15.26}, N_{t-1}, U_t^{27+}$	0.10
$U_{t-2}^{27+}, U_{t-1}^{27+}, U_t^{27+}$	0.08	$U_{t-2}^{27+}, N_{t-1}, U_t^{27+}$	0.07

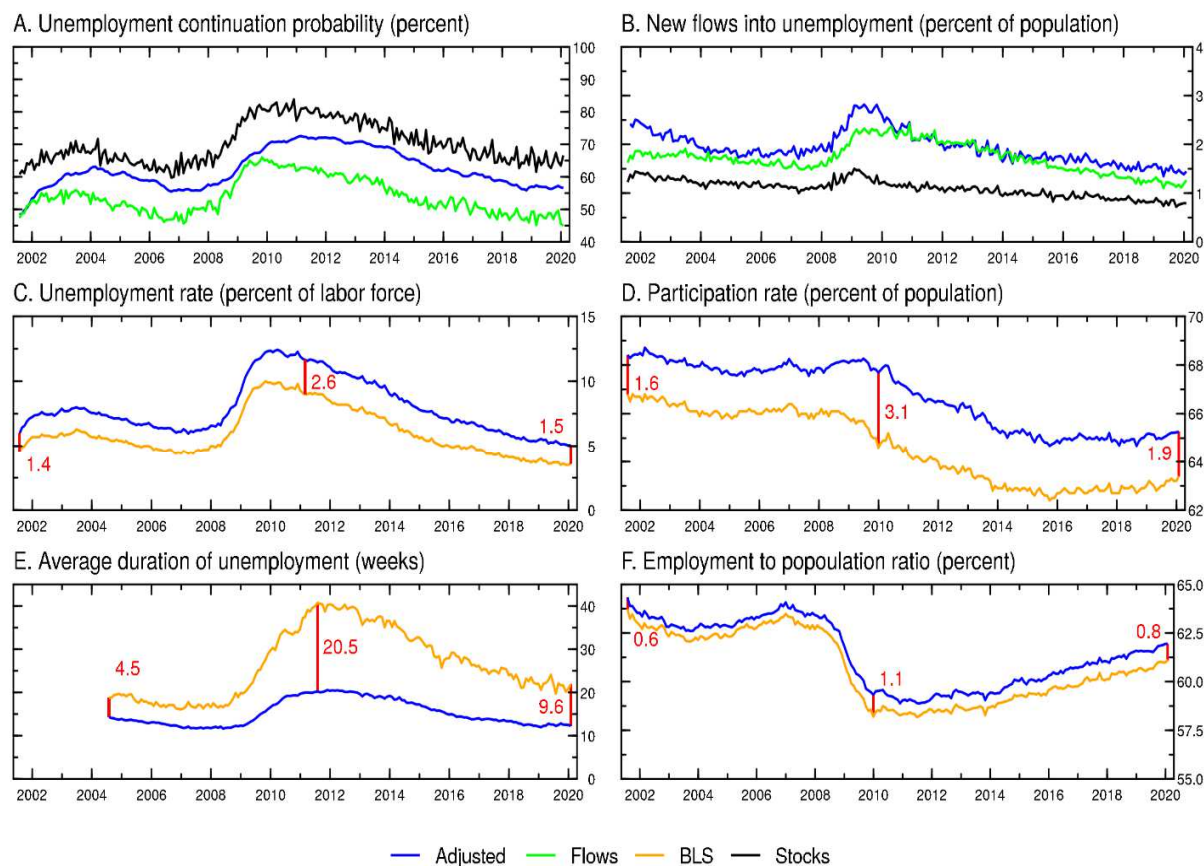
Table 10. Effects of adjustments on unemployment rate and labor-force participation rate.

	Unemployment rate	Labor-force participation rate	Employment-population ratio
Unadjusted BLS	6.0%	64.6%	60.7%
Corrected for rotation-group bias only	6.5%	65.7%	61.4%
Corrected for rotation-group bias and missing observations	6.9%	65.9%	61.4%
Corrected for rotation-group bias, missing observations, and long-term unemployed	8.0%	66.7%	61.4%

Table 11. Adjusted and unadjusted estimates of duration of unemployment


	BLS	Adjusted
< 5 weeks	29.0	36.4
5-14 weeks	27.5	31.6
15-26 weeks	15.4	17.7
> 26 weeks	28.1	14.3
Average duration	26.6 weeks	15.1 weeks

Figure 1. Alternative measures of unemployment-continuation probability, new inflows to unemployment, unemployment rate, labor force participation rate, and average duration of unemployment.




Notes to Figure 1. Panel A: probability that an unemployed individual will still be unemployed next month, Aug 2001 to Feb 2020, as calculated by: (1) ratio of unemployed with duration 5 weeks or greater in month t to total unemployed in $t-1$ (black); (2) fraction of those unemployed in $t-1$ who are still unemployed in t (green); (3) reconciled estimate (blue). Panel B: Number of newly unemployed as a percent of the noninstitutional adult population, Aug 2001 to Feb 2020, as calculated by: (1) number of unemployed with duration less than 5 weeks (black); (2) EU and NU flows as adjusted by BLS (green); (3) reconciled estimate (blue). Panel C: Unemployment rate, July 2001 to Jan 2020, as calculated by BLS (orange) and adjusted estimate (blue). Panel D: labor-force participation rate, July 2001 to Jan 2020, as calculated by BLS (orange) and adjusted estimate (blue). Panel E: Average duration of unemployment, July 2004 to Jan 2020, as calculated by BLS (orange) and adjusted estimate (blue). Panel F: Employment-to-population ratio, July 2001 to Jan 2020, as estimated by BLS (orange) and adjusted (blue). All series seasonally adjusted.

Figure 2. Predicted and reported durations of unemployment for individuals in rotation 1.



Notes to Figure 2. Top panel: reported percentage (blue) and predicted by equation (5) (in yellow) of unemployed who have been searching for indicated number of weeks. Bottom panel: total fraction of unemployed (in black) who have been looking for work for τ weeks and fraction for each type.

Figure 3. Predicted and reported unemployment durations in rotation 2 for individuals who were not in the labor force in rotation 1 and unemployed in rotation 2.



Notes to Figure 3. Horizontal axis: duration of unemployment spell in weeks. Vertical axis: of the individuals who were not in the labor force in rotation 1 and unemployed in rotation 2, the percent who reported having been searching for work at the time of rotation 2 for the indicated duration.

Figure 4. Predicted and actual probability that someone with unemployment duration of τ weeks will still be unemployed next month (top panel) and probability $\eta_2(\tau)$ that the individual is type 2 based on reported duration (bottom panel).


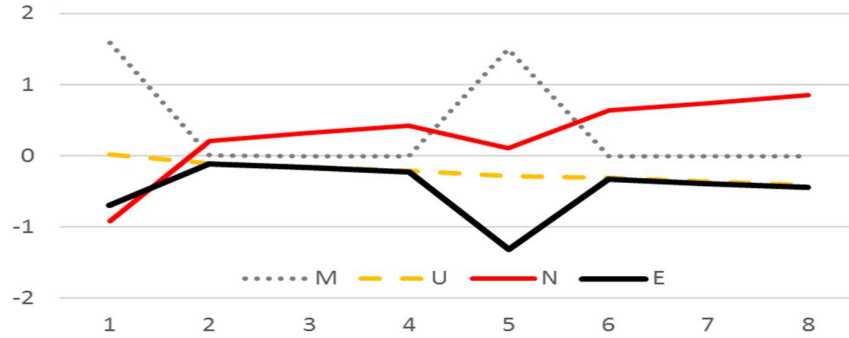


Figure 5. Effect of rotation group on percentage of sampled individuals with indicated reported status.



Notes to Figure 5. Graph shows predicted values implied by regression (13).

Figure 6. Fraction of individuals reporting labor status *E*, *N*, *M*, or *U* in each rotation group (solid blue) and fraction predicted to report that status for that rotation according to equation (25) (dashed red).

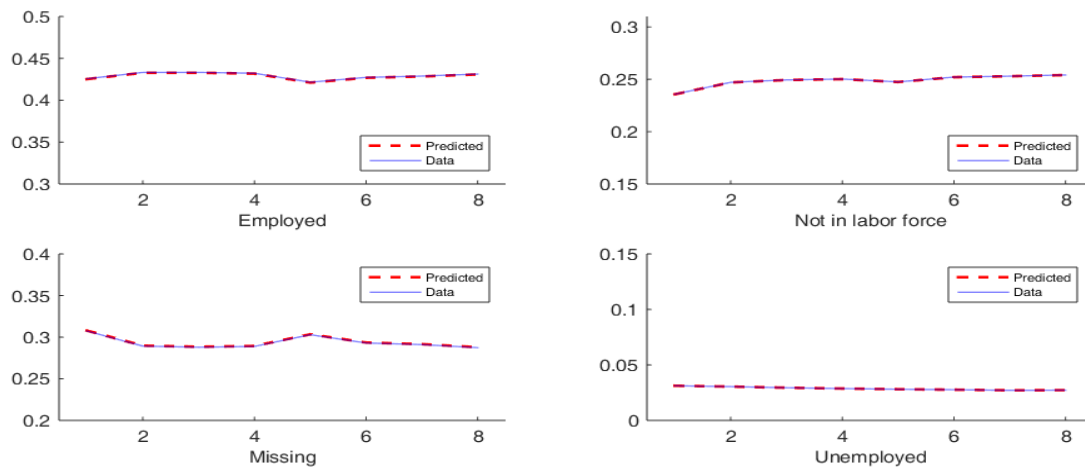


Figure 7. Actual reported transition probabilities for each rotation (solid blue) and fraction predicted by equation (26) (dashed red).

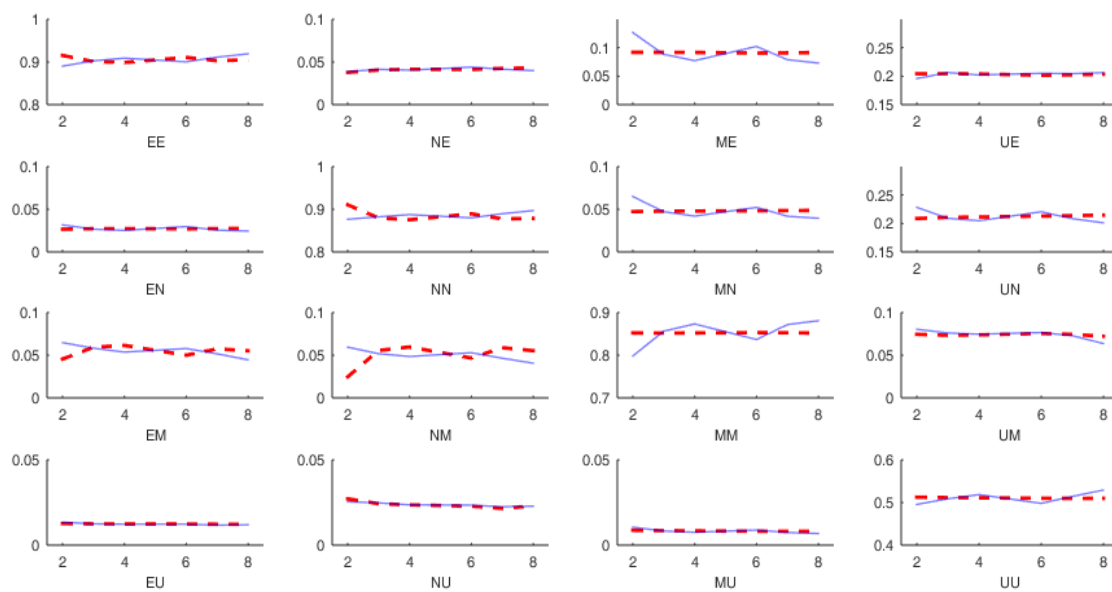


Figure 8. Changes in rotation-group bias parameters over time.


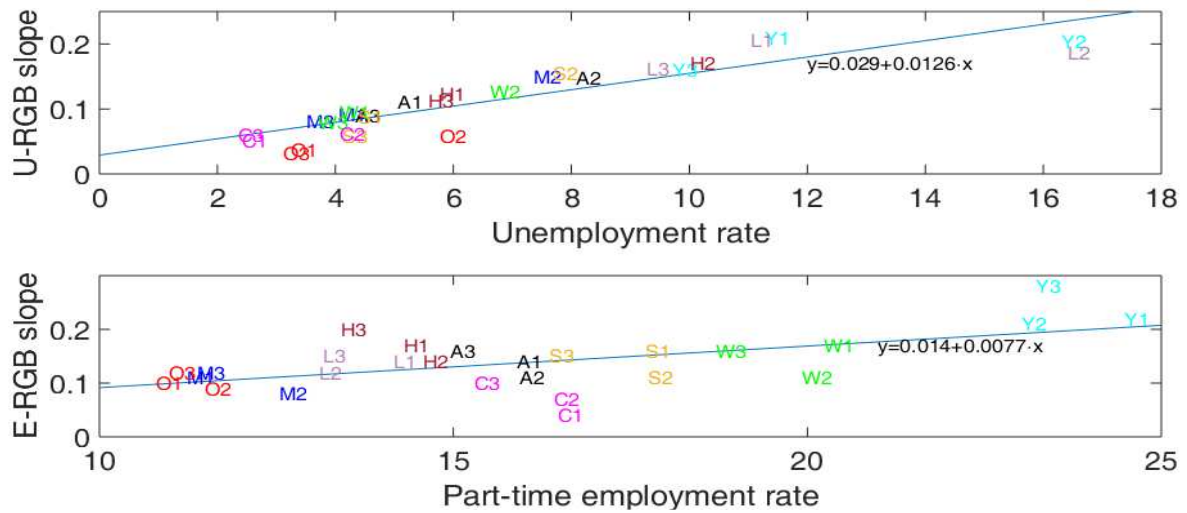



Figure 9. Measures of rotation bias, unemployment, and part-time employment for overall population and different demographic groups over three different subsamples.



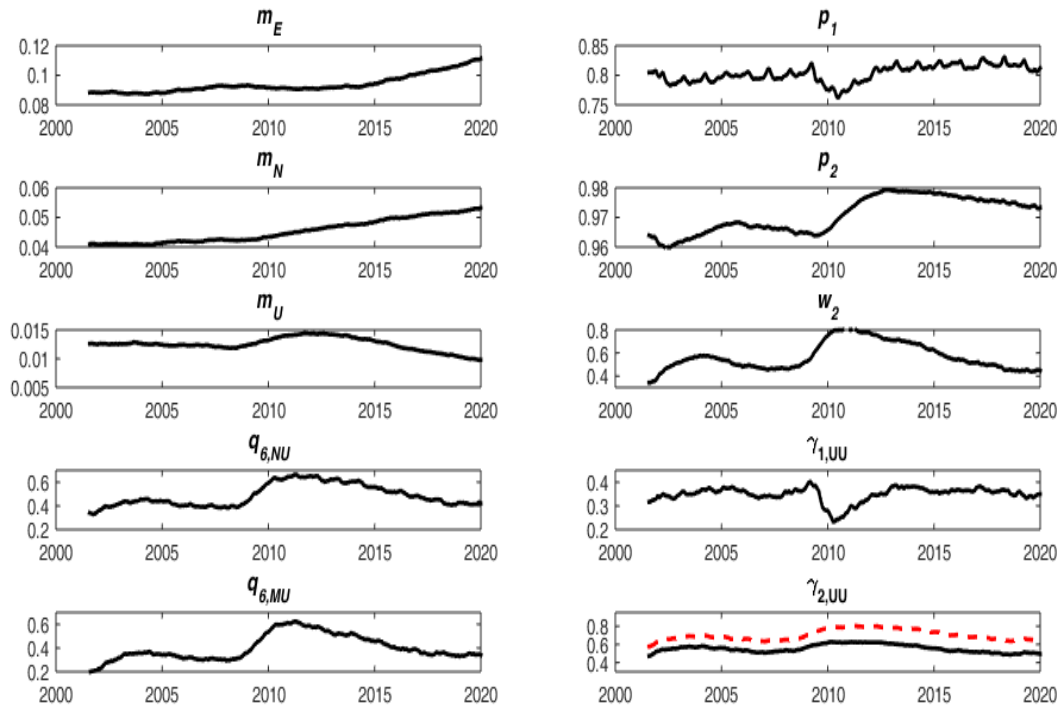
Notes to Figure 9. Top panel: horizontal axis: $U/(U + E)$ as reported by BLS; vertical axis: negative of slope of regression of $U^{[j]}/(U^j + E^{[j]})$ on month-in-sample j . Demographic groups: A: Total population, Y: aged 16-24; M: men aged 25-54; W: women aged 25-54; O: aged 55 and over; L: less than high school education, H: high school graduate; S: some college; C: college graduates. Sample 1: 2001:7-2008:6; sample 2: 2008:7-2014:6; sample 3: 2014:7-2020:1. Also shown is regression line fitted to the 27 observations. Bottom panel: horizontal axis: $\dot{E}/(U + E + N)$ as reported by BLS for \dot{E} part-time employment plus usual full-time workers not at work this month; vertical axis: negative of slope of regression of $\dot{E}^{[j]}/(U^j + E^{[j]} + N^{[j]})$ on month-in-sample j .

Figure 10. BLS reported medium- and long-term unemployment rate (u^{5+}) and percentage of individuals reporting N followed by U^{5+} ($m_N^\#$) for different demographic groups and three different samples plus predicted and actual aggregate values each month.



Notes to Figure 10. Top panel: horizontal axis: $U^{5+}/(U + E)$ as reported by BLS; vertical axis: percentage of population reporting N followed by U^{5+} . Demographic groups: A: Total population, Y: aged 16-24; M: men aged 25-54; W: women aged 25-54; O: aged 55 and over; L: less than high school education, H: high school graduate; S: some college; C: college graduates. Sample 1: 2001:7-2008:6; sample 2: 2008:7-2014:6; sample 3: 2014:7-2020:1. Also shown is regression line fitted to the 27 observations. Bottom panel: month (2001:7 to 2020:1); vertical axis: actual percentage of population reporting N followed by U^{5+} and value predicted from applying the coefficients from demographic regression in the top panel to the aggregate value of $U^{5+}/(U + E)$ the previous month.

Figure 11. Time variation in selected parameters.



Notes to Figure 11. Black lines denote smoothed data summaries $\bar{\theta}_t$ and red dashed lines denote estimates $\tilde{\theta}_t$ that adjust for rotation-group bias, missing observations, and long-term unemployed.

Figure 12. Contributions of different adjustments to the labor-force participation and unemployment rates.





Figure 13. Alternative measures of unemployment rate and new inflows into unemployment.



Notes to Figure 13. Top panel: adjusted unemployment rate (\tilde{u}_t , in blue), BLS unemployment rate (black), U5 (red) and U6 (green). Bottom panel: number of newly unemployed as a percent of the noninstitutional civilian population 16 years and over. Blue: estimate \tilde{V}_t incorporating all adjustments; red: number of unemployed with duration less than 5 weeks; turquoise: number of EU and NU transitions as a fraction of individuals with two consecutive non-missing observations; green: latter adjusted for rotation-group bias alone; black: BLS adjusted EU and NU flows.