

Misclassification Errors and the Underestimation of the U.S. Unemployment Rate *

Shuaizhang Feng[†] Yingyao Hu[‡]

January 28, 2012

Abstract

Using recent results in the measurement error literature, we show that the official U.S. unemployment rate substantially underestimates the true level of unemployment, due to misclassification errors in the labor force status in the Current Population Survey. During the period from January 1996 to August 2011, the corrected monthly unemployment rates are between 1 and 4.4 percentage points (2.1 percentage points on average) higher than the official rates, and are more sensitive to changes in business cycles. The labor force participation rates, however, are not affected by this correction.

Keywords: Unemployment Rate, Labor Force Participation Rate, Misclassification, Measurement Error, Current Population Surveys.

JEL Classification: J21, J64, C14.

*We are grateful to Robert Moffitt and Katheryn Russ for suggestions and comments. We also thank Xianqiang Zou and Jingliang Lu for able research assistance. All remaining errors are our own.

[†]School of Economics, Shanghai University of Finance and Economics, 715 Econ Building, 111 Wuchuan Road, Shanghai 200433, China and IZA. E-mail: shuaizhang.feng@gmail.com

[‡]Corresponding author. Department of Economics, Johns Hopkins University, 440 Mergenthaler Hall, 3400 N. Charles Street, Baltimore, MD 21218, USA and IZA. E-mail: yhu@jhu.edu

1 Introduction

The unemployment rate is among the most important and carefully-watched economic indicators in modern society, and often takes center stage in discussions of economic policy. However, there is considerable disagreement over the precise definition and measurement of unemployment, hence the other two labor force statuses: “employed” and “not-in-labor-force”.¹ In the United States, the Bureau of Labor Statistics (BLS) reports six alternative measures of unemployment (U1-U6), including the official unemployment rate (U3) which is based on the International Labor Organization (ILO)’s definition.² Due to the intrinsic difficulties in classifying some groups of people, such as marginally-attached workers and involuntary part-time workers, into three distinct labor force statuses, the U.S. official unemployment rate is potentially subject to measurement error.

In this paper, we take a latent variables approach and view the reported labor force statuses as functions of the underlying unobserved true labor force statuses. We then impose a structure on the misclassification process and the dynamics of the underlying latent LFS. Using recent results in the measurement error literature, we show that the official U.S. unemployment rate substantially underestimates the true level of unemployment. During the period from January 1996 to August 2011, our corrected unemployment rates are higher than the corresponding official figures by 2.1 percentage points on average. In terms of the monthly differences, the corrected rates are from 1 to 4.4 percentage points higher than the official rates, and are more sensitive to changes in the business cycles.

Official unemployment statistics in U.S. are based on the Current Population Survey (CPS) conducted by the Census Bureau. The CPS interviews around 60,000 households each month to collect basic demographic and labor force status information.

¹For example, using Canadian data, Jones and Riddell (1999) empirically examine labor market transitions of people with different labor force statuses and find substantial heterogeneity within the nonemployed, such that no dichotomy exists between those unemployed and not-in-labor-force among all nonemployed persons.

²The ILO defines “unemployed” as those who are currently not working but are willing and able to work for pay, currently available to work, and have actively searched for work.

Based on the answers to survey questions on job-related activities, the CPS records each individual's labor force status as "employed", "unemployed," or "not-in-labor-force." The misclassification among the three possible values of the labor force status has been a substantial issue in the CPS, as clearly demonstrated by the Reinterview Surveys, in which a small sub-sample of the households included in the original CPS are recontacted and asked the same questions. Treating the CPS reconciled Reinterview Surveys sample as reflecting true labor force statuses, researchers have found that there exists considerable error in the original CPS.³ Of course, the actual misclassification errors in labor force status are likely to be substantially larger than suggested in reconciled CPS reinterviews, as argued by Poterba and Summers (1995), Biemer and Forsman (1992) and Sinclair and Gastwirth (1996).

The misclassification of labor force statuses in the CPS and other similar surveys has received considerable attention in the literature. To identify the misclassification probabilities, early studies typically relied on a particular exogenous sources of "truth", such as the reconciled CPS reinterview surveys (see e.g. Abowd and Zellner 1985, Poterba and Summers 1986, and Magnac and Visser 1999). Nevertheless, the reinterview surveys are usually small in sample size (approximately 3% of the corresponding CPS sample) and not readily available to outside researchers. Reinterview surveys are also subject to misclassification errors due to many practical limitations.⁴ Actually, some studies using other methods show that the reconciled CPS reinterview data may contain even more errors than the original CPS sample (Sinclair and Gastwirth, 1996). Other studies rely on two repeated measures of the labor force status of the same individuals in the same period and assume that the error probabilities are the same for different sub-samples.⁵ More recent studies, such as Biemer and Bushery

³The CPS reinterview sample consists two components, one is "non-reconciled", in which case no attempt is made to determine which answers are correct, the other is "reconciled", in which case the second interviewer would compare the responses from the first survey with the reinterview answers and try to resolve any conflicts (Poterba and Summers, 1984).

⁴The reinterview may not have been independent of the original interview to the extent that respondents remembered and repeated their answers from the original interview. In addition, several factors make it difficult to conduct the reinterview as an exact replication of the original interview, including (1) Only senior interviewers conducted the reinterview, (2) Almost all reinterviews were conducted by telephone, even if the original interview was conducted in person, and (3) The reinterview may not perfectly "anchor" respondents in the original interview's reference period.

⁵See Sinclair and Gastwirth (1996, 1998), which use the H-W model first proposed by Hui and

(2000) and Bassi and Trivellato (2008), explore the panel nature of the surveys and treat the underlying true labor force status as a latent process to be jointly modeled with the misclassification process.

Most existing studies focus on adjusting flows, i.e., the gross labor flows between two consecutive months, not stocks, such as the unemployment rate and the labor market participation rate. While those studies acknowledge that misclassification errors cause serious problems for flows, they somewhat surprisingly assume that errors tend to cancel out for stocks (e.g. Singh and Rao 1995). The only study that has tried to correct for the unemployment rate is Sinclair and Gastwirth (1998). However, their results rely on a key identification assumption that males and females have the same misclassification error probabilities, which we reject in this paper.

This paper uses recent results in the measurement error literature to identify the misclassification probabilities (Hu, 2008). Our method relies only on short panels formed by matching the CPS monthly data sets, thus avoiding the use of auxiliary information such as the reinterview surveys, which are usually small and subject to errors. Our approach is close to the Markov Latent Class Analysis (MLCA) method proposed by Biemer and Bushery (2000), but we use an eigenvalue-eigenvector decomposition to establish a closed-form global identification, while they took a maximum likelihood approach with local identifiability. Generally speaking, parametric GMM or MLE methods typically rely on a local identification argument that the number of unknowns does not exceed that of the restrictions. Given the observed distribution, our identification and estimation procedure directly leads to the unique true values of the unknown probabilities without using the regular optimization algorithms. Therefore, we do not need to be concerned about choosing initial values or obtaining a local maximum in the estimation procedure. In that sense, our estimates are more reliable than those based on local identification, including Biemer and Bushery (2000). Our assumption regarding the dynamics of the underlying true labor force status is also weaker than their first-order Markov chain assumption. In addition, Biemer and Bushery (2000) use group-level data, which are subject to potential biases from within-group heterogeneities. Our identification results enable us to take advantage

Walter (1980).

of the large sample size of the individual-level CPS data, and therefore, to achieve more efficient estimates.

To control for individual heterogeneities, we separately estimate the misclassification probabilities for each demographic group, defined by individual's gender, race and age. Based on those misclassification probabilities, we then estimate the corrected monthly unemployment rates and the labor force participation rates for all demographic groups, and for the US population as a whole. During the period from January 1996 to August 2011, our corrected unemployment rates are higher than the official ones by up to 4.4 percentage points and on average by 2.1 percentage points, with the differences always statistically significant. The most substantial misclassification errors occur when unemployed individuals misreport as either not-in-labor-force or employed. On the other hand, the corrected labor force participation rates and the official ones are rather close and never statistically significantly different.

The rest of the paper is organized as follows. Section 2 provides theoretical results on the identification and estimation of the misclassification probabilities and the marginal distribution of the underlying labor force status. Section 3 presents our main empirical results on the estimated misclassification probabilities and the corrected unemployment rates, along with reported (official) ones. The last section concludes. Additional estimates and simulation results are included in the online appendix of the paper.

2 A closed-form identification result

This section presents a closed-form identification and estimation procedure, which uniquely maps the directly estimable distribution of the self-reported labor force status to the misclassification probabilities and the distribution of the underlying true labor force status. We also evaluate the validity and robustness of the assumptions made in order to achieve identification.

2.1 Assumptions and identification results

Let U_t denote the self-reported labor force status in month t , and X be a vector of demographic variables such as gender, race and age. By matching the monthly CPS samples, we observe the self-reported labor force status in three periods $(t+1, t, t-9)$, together with the demographic variables X for each individual i .⁶ For example, if U_t stands for the labor force status of an individual in January 2008, then U_{t+1} and U_{t-9} denote his or her labor force status in February 2008 and in April 2007, respectively. We denote the i.i.d. sample as $\{U_{t+1}, U_t, U_{t-9}, X\}_i$ for $i = 1, 2, \dots, N$. The self-reported labor force status U_t is defined as follows:

$$U_t = \begin{cases} 1 & \text{employed} \\ 2 & \text{unemployed} \\ 3 & \text{not-in-labor-force} \end{cases}.$$

We denote the latent true labor force status at period t as U_t^* , which takes the same possible values as U_t . Let $\Pr(\cdot)$ stand for the probability distribution function of its argument, we outline our assumptions as follows.

Assumption 1. *The distribution of misclassification errors only depends the true labor force status in the current period, conditional on individual characteristics, i.e.,*

$$\Pr(U_t | U_t^*, X, \mathcal{U}_{\neq t}) = \Pr(U_t | U_t^*, X)$$

for all t with $\mathcal{U}_{\neq t} = \{(U_\tau, U_\tau^*), \text{ for } \tau \neq t\}$.

Assumption 1 still allows the misclassification errors to be correlated with the true labor force status U_t^* and other variables in other periods through U_t^* . This is weaker

⁶Our identification strategy requires matching of three CPS monthly data sets in order to identify the misclassification matrix for the month in the middle of the three months. We choose one month later, i.e., $t+1$, and nine month earlier, i.e., $t-9$, for the following reasons: 1) we want the three periods to be close enough to minimize attrition in CPS samples; 2) we want the three months to cover the 8-month recess period in the CPS rotation structure so that there are enough variations in the labor force status; 3) Assumption 2 on the dynamics of the latent true labor force status is more likely to be satisfied if we use the data reported a while ago, e.g., nine months earlier.

than the classical measurement error assumption, where the error is independent of everything else, including the true values. Assumption 1 is a standard assumption in the literature and allows the misreporting behavior to be summarized by a simple misclassification matrix. Moreover, Meyer (1988) examines this assumption and finds it likely to be valid for CPS data. Assumption 1 implies that the joint probability of the observed labor force status $\Pr(U_{t+1}, U_t, U_{t-9}|X)$ is associated with the unobserved ones as follows:

$$\begin{aligned} & \Pr(U_{t+1}, U_t, U_{t-9}|X) \\ = & \sum_{U_{t+1}^*} \sum_{U_t^*} \sum_{U_{t-9}^*} \Pr(U_{t+1}|U_{t+1}^*, X) \Pr(U_t|U_t^*, X) \Pr(U_{t-9}|U_{t-9}^*, X) \Pr(U_{t+1}^*, U_t^*, U_{t-9}^*|X). \end{aligned} \quad (1)$$

Having established the conditional independence of the misclassification process, our next assumption deals with the dynamics of the latent true labor force status.

Assumption 2. *Conditional on individual characteristics, the true labor force status nine months ago has no predictive power over the true labor force status in the next period beyond the current true labor force status, i.e.,*

$$\Pr(U_{t+1}^*|U_t^*, U_{t-9}^*, X) = \Pr(U_{t+1}^*|U_t^*, X)$$

for all t .

Biemer and Bushery (2000) impose a first-order Markov restriction on the dynamics of the latent labor force status, which states $\Pr(U_{t+1}^*|U_t^*, U_{t-1}^*, \dots, U_1^*) = \Pr(U_{t+1}^*|U_t^*)$. That assumption is likely to be too strong due to the presence of state dependency, serial correlation among idiosyncratic shocks, and unobserved heterogeneity (see e.g. Hyslop 1999). Our assumption 2 is considerably weaker because we use the true labor force status nine month ago. Under Assumption 2, equation (1) may be simplified as follows:

$$\begin{aligned} & \Pr(U_{t+1}, U_t, U_{t-9}|X) \\ = & \sum_{U_t^*} \Pr(U_{t+1}|U_t^*, X) \Pr(U_t|U_t^*, X) \Pr(U_t^*, U_{t-9}|X). \end{aligned} \quad (2)$$

Following the identification results in Hu (2008), we show that all the probabilities containing the latent true labor force status U_t^* on the right-hand-side (RHS) of Equation (2) may be identified under reasonable assumptions. Integrating out U_{t+1} in Equation (2) leads to

$$\Pr(U_t, U_{t-9}|X) = \sum_{U_t^*} \Pr(U_t|U_t^*, X) \Pr(U_t^*, U_{t-9}|X). \quad (3)$$

Following Hu (2008), we introduce our matrix notation. For any given subpopulation with individual characteristics $X = x$, we define the misclassification matrix as follows.

$$\begin{aligned} & M_{U_t|U_t^*, x} \\ \equiv & \begin{bmatrix} \Pr(U_t = 1|U_t^* = 1, x) & \Pr(U_t = 1|U_t^* = 2, x) & \Pr(U_t = 1|U_t^* = 3, x) \\ \Pr(U_t = 2|U_t^* = 1, x) & \Pr(U_t = 2|U_t^* = 2, x) & \Pr(U_t = 2|U_t^* = 3, x) \\ \Pr(U_t = 3|U_t^* = 1, x) & \Pr(U_t = 3|U_t^* = 2, x) & \Pr(U_t = 3|U_t^* = 3, x) \end{bmatrix} \\ \equiv & [\Pr(U_t = i|U_t^* = k, X = x)]_{i,k}. \end{aligned}$$

Each column of the matrix $M_{U_t|U_t^*, x}$ describes how an individual (mis)reports his or her labor force status given a possible value of the true labor force status. The matrix $M_{U_t|U_t^*, x}$ contains the same information as the misclassification probabilities $\Pr(U_t|U_t^*, x)$, which means the identification of $M_{U_t|U_t^*, x}$ implies that of $\Pr(U_t|U_t^*, x)$. Similarly, we may define

$$\begin{aligned} M_{U_t, U_{t-9}|x} & \equiv [\Pr(U_t = i, U_{t-9} = k|x)]_{i,k}, \\ M_{U_t^*, U_{t-9}|x} & \equiv [\Pr(U_t^* = i, U_{t-9} = k|x)]_{i,k}, \\ M_{1, U_t, U_{t-9}|x} & \equiv [\Pr(U_{t+1} = 1, U_t = i, U_{t-9} = k|x)]_{i,k}. \end{aligned}$$

We also define a diagonal matrix as follows:

$$D_{1|U_t^*,x} \equiv \begin{bmatrix} \Pr(U_{t+1} = 1|U_t^* = 1, x) & 0 & 0 \\ 0 & \Pr(U_{t+1} = 1|U_t^* = 2, x) & 0 \\ 0 & 0 & \Pr(U_{t+1} = 1|U_t^* = 3, x) \end{bmatrix}.$$

As shown in Hu (2008), Equations (2) and (3) imply the following two matrix equations:

$$M_{1,U_t,U_{t-9}|x} = M_{U_t|U_t^*,x} D_{1|U_t^*,x} M_{U_t^*,U_{t-9}|x} \quad (4)$$

and

$$M_{U_t,U_{t-9}|x} = M_{U_t|U_t^*,x} M_{U_t^*,U_{t-9}|x}. \quad (5)$$

In order to solve for the unknown matrix $M_{U_t|U_t^*,x}$, we need a technical assumption as follows.

Assumption 3. *The distributions of the current self-reported labor force status conditional on different self-reported labor force statuses nine month ago are linearly independent, i.e., $\Pr(U_t|U_{t-9} = 1, x)$ is not equal to a linear combination of $\Pr(U_t|U_{t-9} = 2, x)$ and $\Pr(U_t|U_{t-9} = 3, x)$ for all U_t and x .*

This assumption is equivalent to the condition that the matrix $M_{U_t,U_{t-9}|x}$ is invertible. Since it is imposed directly on the observed probabilities, this assumption is directly testable. Under Assumption 3, Equation (5) implies that both $M_{U_t|U_t^*,x}$ and $M_{U_t^*,U_{t-9}|x}$ are invertible. Eliminating matrix $M_{U_t^*,U_{t-9}|x}$ in Equations (4) and (5) leads to

$$M_{1,U_t,U_{t-9}|x} M_{U_t,U_{t-9}|x}^{-1} = M_{U_t|U_t^*,x} D_{1|U_t^*,x} M_{U_t|U_t^*,x}^{-1}. \quad (6)$$

This equation implies that the observed matrix on the left-hand-side (LHS) has an eigenvalue-eigenvector decomposition on the RHS. The three eigenvalues are the three diagonal entries in $D_{1|U_t^*,x}$ and the three corresponding eigenvectors are the three columns in $M_{U_t|U_t^*,x}$. Note that each column of $M_{U_t|U_t^*,x}$ is a distribution so that the column sum is 1, which implies that the eigenvectors are normalized.

In order to make the eigenvector unique for each given eigenvalue, we need the eigenvalues to be distinctive, which is formally stated as follows.

Assumption 4. *A different true labor force status leads to a different probability of reporting “employed” in the next period, i.e., $\Pr(U_{t+1} = 1|U_t^* = k, x)$ are different for different $k \in \{1, 2, 3\}$.*

This assumption is also testable from Equation (6). This is because $\Pr(U_{t+1} = 1|U_t^* = k, x)$ for $k \in \{1, 2, 3\}$ are eigenvalues of the observed matrix $M_{1,U_t,U_{t-9}|x}M_{U_t,U_{t-9}|x}^{-1}$. Therefore, Assumption 4 holds if and only if all the eigenvalues of $M_{1,U_t,U_{t-9}|x}M_{U_t,U_{t-9}|x}^{-1}$ in Equation (6) are distinct. Intuitively, this assumption implies that the true labor force status at period t has an impact on the probability of reporting to be employed one period later.

The distinct eigenvalues guarantee the uniqueness of the eigenvectors. Since we do not observe U_t^* in the sample, we need to reveal the value u_t^* for each eigenvector $\Pr(U_t|U_t^* = u_t^*, x)$. In other words, the ordering of the eigenvalues or the eigenvectors is still arbitrary in Equation (6). In order to eliminate this ambiguity, we make the following assumption.

Assumption 5. *Each individual is more likely to report the true labor force status than to report any other possible values, i.e.,*

$$\Pr(U_t = k|U_t^* = k, x) > \Pr(U_t = j|U_t^* = k, x) \text{ for } j \neq k.$$

This assumption does not reveal the value of these misclassification probabilities, nor require the probability of reporting the truth to be larger than 50%. Assumption 5 is consistent with results from CPS reinterviews (see e.g.: Poterba and Summers, 1984) and other validation studies discussed in Bound et al. (2001).

Technically, Assumption 5 implies that the true labor force status is the mode of the conditional distribution of the self-reported labor force status in each column of the eigenvector matrix. Therefore, the ordering of the eigenvectors is fixed and the eigenvector matrix $M_{U_t|U_t^*,x}$ is uniquely determined from the eigenvalue-eigenvector decomposition of the observed matrix $M_{1,U_t,U_{t-9}|x}M_{U_t,U_{t-9}|x}^{-1}$. In particular, after diagonalizing the directly-estimable matrix $M_{1,U_t,U_{t-9}|x}M_{U_t,U_{t-9}|x}^{-1}$, we rearrange the order of

the eigenvectors such that the largest element of each column or each eigenvector, i.e., the mode of the corresponding distribution, is on the diagonal of the eigenvector matrix. Consequently, the misclassification probability $\Pr(U_t|U_t^*, X)$ may be expressed as a closed-form function of the observed probability $\Pr(U_{t+1}, U_t, U_{t-9}|X)$. Such a procedure is constructive because one may estimate the misclassification probability $\Pr(U_t|U_t^*, X)$ by following the identification procedure above.

We summarize the closed-form identification and estimation of the misclassification probability $\Pr(U_t|U_t^*, X)$ as follows.

Theorem 1. *Under Assumptions 1, 2, 3, 4, and 5, the misclassification matrix $\Pr(U_t|U_t^*, X)$ is uniquely determined by the observed joint probability of the self-reported labor force status in three periods, i.e., $\Pr(U_{t+1}, U_t, U_{t-9}|X)$, through the unique eigenvalue-eigenvector decomposition in equation (6).*

Proof: The results directly follow from Theorem 1 in Hu (2008). A complete proof can be found in the online appendix.

Finally, we may estimate the distribution of the latent true labor force status $\Pr(U_t^*|X)$ using the misclassification probability $\Pr(U_t|U_t^*, X)$ from the following equation:

$$\Pr(U_t|X) = \sum_{U_t^*} \Pr(U_t|U_t^*, X) \Pr(U_t^*|X).$$

This equation implies

$$\begin{bmatrix} \Pr(U_t = 1|x) \\ \Pr(U_t = 2|x) \\ \Pr(U_t = 3|x) \end{bmatrix} = M_{U_t|U_t^*, x} \times \begin{bmatrix} \Pr(U_t^* = 1|x) \\ \Pr(U_t^* = 2|x) \\ \Pr(U_t^* = 3|x) \end{bmatrix}.$$

Since we have identified the misclassification probability $\Pr(U_t|U_t^*, X)$, we may solve for the distribution of the latent true labor force status $\Pr(U_t^*|X)$ from that of the self-reported labor force status $\Pr(U_t|X)$ by inverting the matrix $M_{U_t|U_t^*, x}$. Therefore, the

distribution of the latent true labor force status for a given x is identified as follows:

$$\begin{bmatrix} \Pr(U_t^* = 1|x) \\ \Pr(U_t^* = 2|x) \\ \Pr(U_t^* = 3|x) \end{bmatrix} = M_{U_t|U_t^*,x}^{-1} \times \begin{bmatrix} \Pr(U_t = 1|x) \\ \Pr(U_t = 2|x) \\ \Pr(U_t = 3|x) \end{bmatrix}. \quad (7)$$

Given the marginal distribution of the demographic characteristics X , $\Pr(X)$, we may identify the marginal distribution of the latent true labor force status $\Pr(U_t^*)$ as follows

$$\Pr(U_t^*) = \sum_X \Pr(U_t^*|X) \Pr(X).$$

This gives the unemployment rate

$$\mu_t^* \equiv \frac{\Pr(U_t^* = 2)}{\Pr(U_t^* = 1) + \Pr(U_t^* = 2)},$$

and the labor force participation rate

$$\rho_t^* \equiv \Pr(U_t^* = 1) + \Pr(U_t^* = 2).$$

Our identification procedure is constructive as it leads directly to an estimator. A nice property of our approach is that if there is no misclassification error in the data, our estimator would produce the same unemployment rate and labor force participation rate as those based on the raw data, under the assumptions above. Our estimator does not require an initial consistent estimate or iterations as in the regular optimization algorithms do.

2.2 Evaluation of the assumptions

Before proceeding to empirical work, we evaluate the key assumptions which are essential for our identification results. We perform extensive Monte Carlo simulations to examine the robustness of our estimator to deviations from Assumptions 1 and 2. We also test the validity of Assumptions 3 and 4 directly using CPS data. For

Assumption 5, we argue that it is likely to hold based on previous empirical work in the literature. We summarize the main things we have done here while leaving all detailed results in the online appendix.

Assumption 1 imposes conditional independence of the misreporting process. We have considered three different kinds of deviations to this assumption in our Monte Carlo simulations. In the first case, we allow misreporting errors to be correlated with the latent true labor force status in the previous period, i.e., $\Pr(U_t|U_t^*, \mathcal{U}_{\neq t}) = \Pr(U_t|U_t^*, U_{t-1}^*)$. In the second case, misreporting errors may be correlated with the self-reported labor force status in the previous period, i.e., $\Pr(U_t|U_t^*, \mathcal{U}_{\neq t}) = \Pr(U_t|U_t^*, U_{t-1})$. Lastly, we consider a special case of a general relaxation of Assumption 1, i.e., $\Pr(U_t|U_t^*, \mathcal{U}_{\neq t}) = \Pr(U_t|U_t^*, U_{t-1}^*, U_{t-1})$, where people would report the same value as in the previous period with certain probability if their true labor force status does not change, otherwise, they would report following the baseline misclassification probability $\Pr(U_t|U_t^*)$.⁷ In all cases, our simulation results show that our estimator is robust to reasonable deviations from Assumption 1.⁸

Similarly, Assumption 2 imposes conditional independence on the transition of the underlying true labor force status. In the Monte Carlo simulation setup, we relax this assumption to allow the transition of the true LFS to depend on that 9 periods earlier, i.e., $\Pr(U_{t+1}^*|U_t^*, U_{t-9}^*) \neq \Pr(U_{t+1}^*|U_t^*)$. Our simulation results show that the estimator is robust to reasonable deviations to assumption 2.⁹

Assumption 3 requires an observed matrix to be invertible, and therefore, is directly testable from the CPS data. We use bootstrapping to show that the determinant of this matrix is significantly different from zero, which implies that the matrix is invertible.¹⁰

Under Assumptions 1, 2, and 3, Assumption 4 requires that the eigenvalues of an

⁷We do this in response to a referee's concern that reporting behaviors might be serially-correlated.

⁸The detailed Monte Carlo setup can be found at section 3.1.2 in the online appendix and the simulation results can be found at sections 3.2.2-3.2.4 in the online appendix.

⁹The detailed Monte Carlo setup can be found at section 3.1.3 in the online appendix and the simulation results can be found at section 3.2.5 in the online appendix.

¹⁰Detailed results can be found at section 4 (Table A11) in the online appendix.

observed matrix be distinct. We may also directly test this assumption using the CPS data by estimating the differences between the eigenvalues. Our bootstrapping results show that the absolute differences between the eigenvalues are significantly different from zero, which implies that the eigenvalues are distinctive.¹¹

Assumption 5 implies that each individual is more likely to report the true labor force status than any other possible values. We believe this assumption is intuitively reasonable. Also, we are not aware of any studies in the literature (see e.g. previous studies cited in our paper and those reviewed by Bound et al. (2001)) that report anything in violation of this assumption.

3 Empirical Results

3.1 Matching of monthly CPS data

We use the public-use micro CPS data to estimate the unemployment rate and the labor force participation rate.¹² Each CPS monthly file contains eight rotation groups that differ in month-in-sample. The households in each rotation group are interviewed for four consecutive months after they enter, withdraw temporarily for eight months, then reenter for another four months of interviews before exiting the CPS permanently. Because of the rotational group structure, the CPS can be matched to form longitudinal panels, which enable us to obtain the joint probabilities of the self-reported labor force statuses in three periods.

We follow the algorithm proposed by Madrian and Lefgren (2000) to match adjacent CPS monthly files.¹³ There are two main steps in the process of matching. First, the CPS samples are matched based on identifiers. If two individuals in two CPS monthly files (within the corresponding rotational groups) have the same household identifier,

¹¹Results can be found at Table A12 of section 4 in the online appendix.

¹²All data are downloaded from www.bls.census.gov/cps_ftp.html. Following BLS practices, we restrict the samples to those aged 16 and over. Sample summary statistics can be found at Table A1 in the online appendix.

¹³See also Feng (2001) and Feng (2008).

household replacement number (which denotes whether this is a replacement of the initial household) and personal identifier (which uniquely identifies a person within a household), then the two individuals are declared as a “crude match”. This step is not perfect and may result in considerable matching errors because there might exist coding errors with respect to those identifiers. Therefore, the second step uses information on sex, age and race to “certify” the crude match. In the matching algorithm we use, if the sex or race reported in the two monthly files corresponding to a crude match are different, or if the age difference is greater than 1 or less than 0, then we discard the match as a false one.

As the previous literature (e.g.: Peracchi and Welch 1995 and Feng 2008) has documented, the matched sample is not representative of the cross-sectional sample in period t due to sample attrition in matching. We use the matching weights to correct for attrition. First, we run a Logit regression for the period t cross-sectional sample, where the dependent variable is either 1 (the observation is matched) or 0 (the observation is not matched), and the independent variables are sex, race, age, schooling, and the labor force status in period t . We next calculate the predicted probabilities of being matched for all the observations in the matched sample. The final matched sample is then weighted using the inverse of the predicted match probabilities. This adjustment procedure ensures the cross-sectional sample and the matched sample have the same marginal distributions on the key individual characteristics for period t .¹⁴

3.2 Misclassification probabilities

For each demographic group, we pool matched samples to estimate the misclassification probabilities.¹⁵ Table 1 reports results for all the eight groups, including (1)

¹⁴Under the assumption that attrition is solely based on observables, our correction method using match weights is consistent. To check for robustness of our procedure we have also tried not using matching weights, i.e., not correcting for attrition in matching, and found similar results in terms of corrected unemployment rates. Details can be found at section 5.5 of the online appendix.

¹⁵To be consistent with the last version of the paper we pool data from January 1996 to December 2009. The estimated misclassification probabilities do not change statistically significantly if we pool all data up to August 2011. Please refer to section 5.3 of the online appendix for details and more

white males aged 40 and younger; (2) white males aged over 40; (3) nonwhite males aged 40 and younger; (4) nonwhite males aged over 40; (5) white females aged 40 and younger; (6) white females aged over 40; (7) nonwhite females aged 40 and younger; (8) nonwhite females aged over 40. There exist some consistent patterns across all the groups. When the actual labor force status is either employed or not-in-labor-force, the probabilities of being misreported to a different labor force status are typically small and never above 6%. The biggest errors come from the unemployed people being misclassified as either not-in-labor-force or employed. Only around 50-70% of unemployed people correctly report their true labor force status. For example, for white males aged 40 and younger, 20% of the unemployed report to be employed, while another 17% of them report as not-in-labor-force. On the other hand, there are considerable heterogeneities among different demographic groups. For example, 10.8% of the unemployed white females aged 40 and younger report as not-in-labor-force, while all other groups have much higher probabilities of reporting to be not-in-labor-force while unemployed.

We also formally test for the differences in the misclassification probabilities between the groups. For example, we consider males vs. females, controlling for race and age categories. We find that employed males are more likely to misreport as unemployed but less likely to misreport as not-in-labor-force than employed females. The differences are always statistically significant at the 5% significance level except for the comparison between nonwhite males aged 40 and younger and nonwhite females aged 40 and younger. When unemployed, the differences are mostly insignificant, with the only exception being that white males aged over 40 are less likely to misreport as being not-in-labor-force compared to white females aged over 40. In addition, when not-in-labor-force, males are more likely to be misclassified as employed.¹⁶

Some previous studies have made strong assumptions regarding between-group misclassification errors. For example, in order to achieve identification, Sinclair and Gastwirth (1998) assume that males and females have the same misclassification er-

elaborate discussions.

¹⁶Comparisons between males and females and other demographic characteristics can be found in Table A13 in the online appendix.

ror probabilities (see also Sinclair and Gastwirth 1996), which we can safely reject.¹⁷ In general, our results suggest that the equality assumptions of misclassification probabilities across different demographic groups, which are essential for identification in the H-W models, are unlikely to hold in reality.

The last two rows of Table 1 report misclassification probabilities and associated standard errors for the overall U.S. population. The results are broadly consistent with those in the existing literature. When we compare our estimates of misclassification probabilities with some of those obtained in the existing literature,¹⁸ we see the same general pattern: the biggest misclassification probabilities happen when unemployed individuals misreport their labor force statuses as either not-in-labor-force ($\Pr(U_t = 3|U_t^* = 2)$) or employed ($\Pr(U_t = 1|U_t^* = 2)$), while the other misclassification probabilities are all small. Our point estimates of $\Pr(U_t = 3|U_t^* = 2)$ and $\Pr(U_t = 1|U_t^* = 2)$ are somewhat higher than many of the existing estimates. But our estimates are well within the 95% confidence intervals in many existing studies because of their large standard errors. Due to our methodological advantages and the large sample size we use, we are able to produce much more precise estimates.

3.3 The unemployment rate

Given the estimated misclassification matrices, we then calculate distribution of the latent true labor force status for each demographic group based on Equation (7). To estimate $\Pr(U_t|X)$, we use all the eight rotation groups in any given CPS monthly file, which subsequently give us the self-reported unemployment rate and the labor force participation rate. Once we have $\Pr(U_t^*|X)$, we can calculate the corrected unemployment rate and the corrected labor force participation rate. In order to be consistent with officially-announced statistics, all numbers are weighted using final weights provided by CPS.¹⁹

¹⁷See the first panel in Table A13 in the online appendix.

¹⁸These estimates can be found in Table A14 in the online appendix.

¹⁹The final weights in the CPS micro data have been adjusted for a composite estimation procedure that BLS uses to produce official labor force statistics (Appendix I in BLS, 2000).

Table 2 presents the results for each demographic group. We divide the study period into three sub-periods based on the US business cycles.²⁰ The first sub-period goes from January 1996 to October 2001, which is roughly the end of the 2001 recession. The second sub-period is from November 2001 to November 2007, corresponding to the expansion period between two recessions (the 2001 recession and the most recent 2007-09 recession). The third sub-period goes from December 2007 to the end of our study period, i.e., Aug 2011, which includes the 2007-2009 recession and its aftermath.

For each demographic group and each sub-period, the corrected unemployment rates are always higher than the reported ones. Note also that for all demographic groups, sub-period 3 posts the highest levels of unemployment, followed by sub-period 2, and then by the first sub-period. This relationship is unchanged using either the reported or the corrected rates. In addition, the degree of underestimation is larger when the level of unemployment is higher. For example, for white males less than 40, in the first sub-period, the corrected unemployment rate is 6.5%, which is higher than the reported unemployment rate by 1.5 percentage points. In the second sub-period, the corrected unemployment rate is 8.2%, which is higher than the reported unemployment rate by 2.1%. The largest differential appears in the latest recession period. In this case, the corrected unemployment rate is 14.5%, which is higher than the reported unemployment rate by 4.4% – a 44% upward adjustment.

We then estimate the unemployment rates and the corresponding standard errors for the US population as a whole, based on the results for all the demographic groups. Based on the last two rows of Table 2, corrected unemployment rates for the US population are 5.9%, 6.9% and 11.5% for the three sub-periods, respectively. Note that the degree of underestimation is substantially larger in the third sub-period, official unemployment rate is 3.4 percentage points lower than the corrected one, while in the first two sub-periods the discrepancies are only 1.5 and 1.8 percentage points, respectively. Figure 1 displays all the monthly values that are seasonally-adjusted. For the whole period, the corrected unemployment rate is always higher than the reported one and the difference is between 1% and 4.4%, and 2.1% on average.

²⁰see <http://www.nber.org/cycles.html>.

The substantial degree of underestimation of the unemployment rate may not be very surprising because most of the misclassification errors are from the unemployed people misreporting their labor force status as either employed or not-in-labor-force. We believe this arises primarily due to the intrinsic difficulties in classifying labor force status for some specific groups of people. Among those not-in-labor-force, marginally-attached workers, especially discouraged workers, could be classified as unemployed because they also desire a job although do not search in the job market. In fact, Jones and Riddell (1999) find that some marginally-attached workers are behaviorally more similar to unemployed than to the rest of those not-in-labor-force. On the other hand, involuntary part-time workers are classified as employed according to the official definition. But many of them could be observationally more similar to unemployed workers.²¹ ²²

Table 3 decomposes the underestimation of unemployment rate. For the period of January 1996 to August 2011, the official statistics underestimate the unemployment rate on average by 2.1 percentage points. The degree of underestimation varies, however, by demographic group. On the one hand, the young nonwhite female group posts the largest level of underestimation, at 5 percentage points. On the other hand, the official statistics only underestimate by 1.3 percentage points for white males over 40. In terms of contributions to the total degree of underestimation (last column of Table 3), white females over 40 declare the largest share of the total (27%), followed by white males 40 and younger (21%). Nonwhite groups contributed relatively little as they account for relatively small portions of the US total population.

One particular concern is whether misclassification behaviors and the resulted corrected unemployment rates would depend on labor market conditions. For example, when the labor market is weak and the pool of unemployed people includes a larger share of job losers and others whose status is unambiguous, then the misreporting of unemployment would tend to be less prevalent. In order to test this hypothe-

²¹For example, Farber (1999) examine displaced workers and find temporary and involuntary part-time jobs are part of the transitional process from unemployment to full-time work.

²²According to the broadest concept of unemployment by BLS, U6, all marginally-attached workers and involuntary part-time workers are counted as unemployed. Our corrected unemployment rate series are substantially lower than U6, as shown by Figure A4 in the online appendix.

sis directly, we have estimated three different misclassification probabilities for each demographic group for the three sub-periods. We do find some evidence that the misclassification probabilities are different in different sub-period corresponding to different labor market conditions. More specifically, sub-period 3 (December 2007 to August 2011), which is characterized by much higher rate of unemployment and presumably much weaker labor market conditions compared to the previous two sub-periods, has lower levels of misclassification in general. Nevertheless, we show that the corrected unemployment series are robust to whether we allow misclassification probabilities to differ in different sub-periods.²³

We have also examined the effect of misclassification on the labor force participation (LFP) rates. For each demographic group for the three sub-periods: January 1996 to October 2001, November 2001 to November 2007, and December 2007 to August 2011, the corrected labor force participation rates are always higher than the reported ones, but the differences are small and not statistically significant. For the US population as a whole, average difference between corrected and official LFP rates is less than 2%, and not statistically significant. For the three sub-periods, the corrected labor force participation rates are 68.1%, 67.3% and 66.8%, respectively. The reported rates are only slightly lower, at 67.1%, 66.2% and 65.2%, respectively.²⁴ Therefore, misclassification errors cause little change to the labor force participation rate. Compared with the number of unemployed people, the total number of people who are in labor force is much larger. Hence any corrections due to misclassification errors will have a relatively small effect.

4 Conclusion

This paper examines misclassification errors in labor force status using CPS data. Similar to previous studies, we show that there exist considerable misclassifications from unemployed to not-in-labor-force and from unemployed to employed. The results at least partly reflect the intrinsic difficulties in classifying labor force statuses of

²³Detailed results can be found at section 5.4 in the online appendix.

²⁴Detailed results are shown in section 7 of the online appendix.

certain groups of people, such as marginally attached workers (especially discouraged workers) and part-time workers for economic reasons, into three distinct categories. We correct for such errors and show that the official U.S. unemployment rate significantly underestimates the true level of unemployment in the United States. For the period from January 1996 to August 2011, our corrected unemployment rates are higher than the reported ones by 2.1 percentage points on average, with differences ranging from 1 to 4.4 percentage points and always statistically significant. In addition, our estimates suggest that unemployment might be much more sensitive to business cycles than previously thought, as the degree of underestimation is larger in magnitude when unemployment rate is higher.


References

- J.M. Abowd and A. Zellner. Estimating gross labor-force flows. *Journal of Business and Economic Statistics*, 3(3):254–283, 1985.
- F. Bassi and U. Trivellato. A latent class approach for estimating gross flows in the presence of correlated classification errors. In P. Lynn, editor, *Methodology of Longitudinal Studies*. Chichester, Wiley, 2008.
- P.P. Biemer and J.M. Bushery. On the validity of markov latent class analysis for estimating classification error in labor force data. *Survey Methodology*, 26(2):139–152, 2000.
- P.P. Biemer and G. Forsman. On the quality of reinterview data with application to the Current Population Survey. *Journal of the American Statistical Association*, 87(420):915–923, 1992.
- BLS. *Current Population Survey: Design and Methodology*. Bureau of Labor Statistics, 2000. Technical Paper 63RV.
- J. Bound, C. Brown, and N. Mathiowetz. Measurement error in survey data. In *Handbook of Econometrics*, volume 5, pages 3705–3843. 2001.

- H.S. Farber. Alternative and part-time employment arrangements as a response to job loss. *Journal of Labor Economics*, 17(S4):142–169, 1999.
- S. Feng. The longitudinal matching of Current Population Surveys: A proposed algorithm. *Journal of Economic and Social Measurement*, 27(1-2):71–91, 2001.
- S. Feng. Longitudinal matching of recent Current Population Surveys: Methods, non-matches and mismatches. *Journal of Economic and Social Measurement*, 33(4):241–252, 2008.
- Y. Hu. Identification and estimation of nonlinear models with misclassification error using instrumental variables: A general solution. *Journal of Econometrics*, 144: 27–61, 2008.
- S.L. Hui and S.D. Walter. Estimating the error rates of diagnostic tests. *Biometrics*, 36:167–171, 1980.
- D.R. Hyslop. State dependence, serial correlation and heterogeneity in intertemporal labor force participation of married women. *Econometrica*, 67(6):1255–1294, 1999.
- S.R.G. Jones and C.W. Riddell. The measurement of unemployment: an empirical approach. *Econometrica*, 67(1):147–162, 1999.
- B. Madrian and L. Lefgren. An approach to longitudinally matching Current Population Survey (CPS) respondents. *Journal of Economic and Social Measurement*, 26:31–62, 2000.
- T. Magnac and M. Visser. Transition models with measurement errors. *Review of Economics and Statistics*, 81(3):466–474, 1999.
- B.D. Meyer. Classification-error models and labor-market dynamics. *Journal of Business and Economic Statistics*, 6(3):385–390, 1988.
- F. Peracchi and F. Welch. How representative are matched cross sections? evidence from the Current Population Survey. *Journal of Econometrics*, 68(1):153–179, 1995.

- J.M. Poterba and L.H. Summers. Response variation in the CPS: caveats for unemployment analysts. *Monthly Labor Review*, pages 37–43, March 1984.
- J.M. Poterba and L.H. Summers. Reporting errors and labor market dynamics. *Econometrica*, 54(6):1319–1338, 1986.
- J.M. Poterba and L.H. Summers. Unemployment benefits and labor market transitions: a multinomial logit model with errors in classification. *Review of Economics and Statistics*, 77(2):207–216, 1995.
- M.D. Sinclair and J.L. Gastwirth. On procedures for evaluating the effectiveness of reinterview survey methods: application to labor force data. *Journal of the American Statistical Association*, 91(435):961–969, 1996.
- M.D. Sinclair and J.L. Gastwirth. Estimates of the errors in classification in the labour force survey and their effect on the reported unemployment rate. *Survey methodology*, 24(2):157–169, 1998.
- A.C. Singh and J.N.K. Rao. On the adjustment of gross flow estimates for classification error with application to data from the Canadian labour force survey. *Journal of the American Statistical Association*, 90(430):478–488, 1995.

Figure 1: Corrected and official (reported) unemployment rates



Note: Figure displays seasonally-adjusted corrected unemployment rates (in solid line) and official unemployment rates (in dashed line) for the whole population from Jan 1996 to Aug 2011. The corresponding thin lines signify 95% upper and lower confidence bounds. For seasonally adjustment, we use Census Bureau's WinX12 software.

Table 1: Misclassification probabilities (%) for different demographic groups

Demographic group	$\Pr(i j) \equiv \Pr(U_t = i U_t^* = j)$					
	$\Pr(2 1)$	$\Pr(3 1)$	$\Pr(1 2)$	$\Pr(3 2)$	$\Pr(1 3)$	$\Pr(2 3)$
(1) Male/White/age $\leq$ 40	0.9 (0.06)	1.3 (0.07)	20.1 (1.28)	17.2 (2.69)	6.0 (0.42)	0.0 (0.39)
(2) Male/White/age > 40	0.4 (0.03)	0.9 (0.05)	16.5 (1.14)	18.8 (2.34)	1.4 (0.07)	0.1 (0.07)
(3) Male/Nonwhite/age $\leq$ 40	1.1 (0.10)	2.2 (0.13)	13.4 (1.21)	18.1 (3.91)	5.0 (0.36)	4.3 (1.26)
(4) Male/Nonwhite/age > 40	0.7 (0.08)	1.5 (0.10)	15.5 (1.81)	22.0 (5.55)	1.2 (0.16)	0.0 (0.12)
(5) Female/White/age $\leq$ 40	0.6 (0.05)	2.1 (0.10)	18.6 (1.59)	10.8 (4.10)	4.4 (0.27)	0.0 (0.08)
(6) Female/White/age>40	0.3 (0.03)	1.4 (0.07)	17.9 (1.46)	28.2 (3.16)	1.0 (0.06)	0.0 (0.01)
(7) Female/Nonwhite/age $\leq$ 40	1.1 (0.09)	2.6 (0.16)	11.8 (1.54)	29.4 (8.24)	2.2 (0.70)	0.0 (0.01)
(8) Female/Nonwhite/age>40	0.4 (0.07)	1.8 (0.11)	13.9 (1.89)	25.0 (5.82)	1.2 (0.09)	0.7 (0.17)
Overall	0.6 (0.02)	1.5 (0.03)	17.3 (0.59)	20.2 (1.39)	2.9 (0.10)	0.2 (0.09)

Note: Bootstrap standard errors based on 500 repetitions are reported in parentheses.

Table 2: Unemployment rates (%) averaged over three sub-periods for different demographic groups

Demographic group	Sub-period 1		Sub-period 2		Sub-period 3	
	(1996/01-2001/10)		(2001/11-2007/11)		(2007/12-2011/8)	
	reported	corrected	reported	corrected	reported	corrected
(1) Male/White/age $\leq$ 40	5.0 (0.2)	6.5 (0.4)	6.1 (0.3)	8.2 (0.5)	10.1 (0.5)	14.5 (0.8)
(2) Male/White/age $>$ 40	2.7 (0.1)	3.4 (0.2)	3.4 (0.2)	4.5 (0.2)	6.3 (0.3)	8.9 (0.5)
(3) Male/Nonwhite/age $\leq$ 40	10.1 (0.5)	11.1 (0.9)	10.8 (0.5)	12.0 (1.1)	16.0 (0.7)	19.3 (1.4)
(4) Male/Nonwhite/age $>$ 40	4.8 (0.2)	6.5 (0.5)	5.8 (0.3)	8.0 (0.6)	9.6 (0.4)	13.9 (1.0)
(5) Female/White/age $\leq$ 40	5.1 (0.2)	6.4 (0.4)	5.8 (0.3)	7.3 (0.5)	8.3 (0.4)	10.9 (0.7)
(6) Female/White/age $>$ 40	2.7 (0.1)	4.4 (0.3)	3.2 (0.1)	5.3 (0.3)	5.4 (0.2)	9.1 (0.5)
(7) Female/Nonwhite/age $\leq$ 40	10.0 (0.5)	14.5 (1.5)	10.3 (0.5)	14.9 (1.6)	13.4 (0.6)	19.8 (2.0)
(8) Female/Nonwhite/age $>$ 40	4.2 (0.2)	5.1 (0.5)	5.2 (0.2)	6.8 (0.6)	7.2 (0.3)	10.0 (0.9)
Overall	4.4 (0.1)	5.9 (0.2)	5.1 (0.1)	6.9 (0.2)	8.1 (0.1)	11.5 (0.3)

Note: Numbers reported in parentheses are bootstrap standard errors based on 500 repetitions.

Table 3: Decomposition of underestimation in unemployment rates by demographic groups

Demographic group	Underestimation in unemployment rate $(a) = \widehat{\mu}_t^* - \mu_t$	Group share in US population (b)	Contribution to underestimation $(c) = (a) \times (b)$	Relative contribution $(d) = \frac{(c)}{\sum(c)}$
(1) Male/White/age $\leq$ 40	2.41	18.24	0.44	20.57
(2) Male/White/age $>$ 40	1.34	21.80	0.29	13.65
(3) Male/Nonwhite/age $\leq$ 40	1.72	4.46	0.08	3.59
(4) Male/Nonwhite/age $>$ 40	2.65	3.73	0.10	4.63
(5) Female/White/age $\leq$ 40	1.68	17.91	0.30	14.08
(6) Female/White/age $>$ 40	2.37	24.20	0.57	26.82
(7) Female/Nonwhite/age $\leq$ 40	5.05	4.99	0.25	11.79
(8) Female/Nonwhite/age $>$ 40	1.76	4.68	0.08	3.86
Total		100.00	2.14	100.00

Note: Table reports averages over the January 1996 to August 2011 period. All numbers are rounded. (a) Underestimation in the unemployment rate (%), which equals the average corrected unemployment rate $\widehat{\mu}_t^*$ minus the average official unemployment rate μ_t ; (b) Population share of the demographic group; (c) Contribution to the total US underestimation in the unemployment rate (%), which equals (a) times (b); (d) Relative contribution to the total underestimation, which equals (c) divided by its column sum.