

Working Papers

RESEARCH DEPARTMENT

WP 21-22

June 2021

<https://doi.org/10.21799/frbp.wp.2021.22>

Measuring Employer-to-Employer Reallocation

Shigeru Fujita

Federal Reserve Bank of Philadelphia Research Department

Giuseppe Moscarini

Yale University and
Federal Reserve Bank of Philadelphia Research Department Visiting Scholar

Fabien Postel-Vinay

University College London

ISSN: 1962-5361

Disclaimer: This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in these papers are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. Philadelphia Fed working papers are free to download at: <https://philadelphiafed.org/research-and-data/publications/working-papers>.

Measuring Employer-to-Employer Reallocation*

Shigeru Fujita[†] Giuseppe Moscarini[‡] Fabien Postel-Vinay[§]

May 2021

Abstract

We revisit the measurement of Employer-to-Employer (EE) transitions in the monthly Current Population Survey. We detect sharp increases in the incidence of missing answers to the relevant question starting in 2007, when the U.S. Census Bureau introduced the Respondent Identification Policy. We show evidence of non-response selection by both observable and unobservable worker characteristics that correlate with EE mobility. We propose a selection model and a procedure to impute missing answers, thus EE transitions. Our imputed EE aggregate series restores a close congruence with the business cycle after 2007, including the COVID-19 recession, and exhibits no downward trend since 2000.

*We thank Anne Polivka at the BLS for her assistance on various technical aspects of the CPS, and Henry Hyatt (discussant), Jim Spletzer, and participants at the 2019 Montreal Workshop on Markets with Frictions, the Federal Reserve Bank of Atlanta's 10th Annual Employment Conference, and the 2019 BOC/BOJ/Philadelphia Fed joint conference on macroeconomics, for comments and discussions. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of Philadelphia or any other person associated with the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

JEL Codes: J63, E24.

[†]Federal Reserve Bank of Philadelphia, Research Department. shigeru.fujita@phil.frb.org.

[‡]Yale University, Department of Economics and Cowles Foundation; NBER. giuseppe.moscarini@yale.edu.

[§]University College London, Department of Economics and Institute for Fiscal Studies. f.postel-vinay@ucl.ac.uk.


1 Introduction

The labor market in the US is a tremendously dynamic place. Every month, millions of workers move between employment, unemployment, and out of the labor force. In recent years, increasing attention has been paid to the flow of workers from Employer to Employer (EE), with no intervening jobless spell. A prominent literature, as best exemplified by Burdett and Mortensen (1998) and its empirical applications, as well as by Postel-Vinay and Robin (2002), shows that on-the-job search by, and competition between firms for, employed workers are a natural source of worker bargaining power, and an important determinant of cross-sectional wage dispersion caused by turnover frictions.

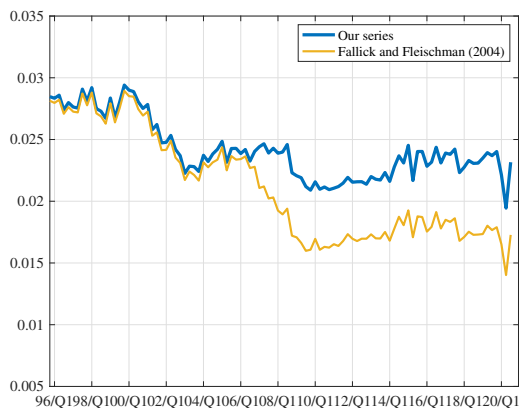
Just as critical is the role that EE reallocation plays in shaping two dynamic aspects of US labor markets. First, from the individual point of view of a typical US worker, direct moves from one employer to another are a major source of earnings growth over the life cycle (Topel and Ward, 1992), but also of idiosyncratic earnings risk. Climbing the job ladder takes time; therefore, falling off it can have drastic implications for lifetime earnings (Davis and Von Wachter, 2011) and explain the striking skewness and kurtosis in individual earnings growth at annual frequency documented by Guvenen et al. (2014) (see, e.g., Hubmer, 2018). Second, from an aggregate point of view, the total EE flow is comparable in size with the flows from Unemployment (UE) and Nonparticipation into Employment (NE). A large share of these UE and NE flows comprises, respectively, recalls by the last employer (Fujita and Moscarini, 2017) and first-time entry into the labor force, which do not directly reallocate workers between firms. Therefore, EE transitions play a quantitatively dominant role in this type of reallocation, which is a major driver of aggregate productivity growth (e.g. Foster et al., 2008 and Lentz and Mortensen, 2008). The EE transition probability is also procyclical, but much less volatile than the UE probability or the unemployment rate. These facts bear significant implications for the cyclical reallocation of labor input between firms, industries, and occupations (Haltiwanger et al., 2018), for the estimation of the matching function (Moscarini and Postel-Vinay, 2018), and for measurement of mismatch and labor market slack relevant to monetary policy (Moscarini and Postel-Vinay, 2019).

For all these reasons — and possibly more — measuring EE transitions accurately is important. This is the goal of the present paper.¹ We focus on the monthly Current Population

¹The type of transitions we focus on involves a change of *employer* — hence the systematic reference to “Employer-to-Employer (EE) transitions.” In the literature, these are sometimes referred to as “Job-to-Job” (J2J): we find this label confusing as, strictly speaking, job changes include internal promotions, demotions, or moves caused by internal restructuring and reorganizations, which typically do not involve a change of employer. We exclude those within-employer job changes from our analysis, although we hereby acknowledge that they are potentially just as relevant to reallocation and productivity growth as EE transitions.



(a) 12-month trailing moving average



(b) quarterly average

Figure 1: Employer-to-Employer (EE) transition probability (Sep 1996 - November 2020)

Survey (CPS), the premier source of real-time information on labor markets, including the civilian unemployment rate, available to policymakers in the United States. The monthly frequency, almost unique even among labor force surveys in developed countries, reduces the recall bias and time aggregation that blur the distinction between direct EE transitions and short unemployment spells in survey data. Since its 1994 redesign, the CPS contains an explicit retrospective question (variable IODP1) whose yes/no answer can be used to identify EE transitions: the interviewer reads out the name of an individual’s employer recorded in the previous month, and asks if it is still the same. In this paper, we will refer to this question as “SAMEMP.” Fallick and Fleischman (2004) pioneered its use to estimate the average EE monthly transition probability, and a time series that has become the standard reference in the profession. The lighter (yellow) line in Figure 1a shows the time series of our replication of their results, after taking a 12-month trailing Moving Average to eliminate high-frequency noise; Figure 1b shows quarterly averages of the seasonally adjusted series. We can see a dramatic decline that starts in early 2007, and never reverts, thus generating the impression of a strong cyclical drop preceding the Great Recession by a full year, as well as a downward trend, and a similarly dramatic but transient drop in April-May 2020 during the COVID-19 crisis.

In this paper, we revisit measurement of the EE transition probability. Our starting point is Figure 2. We detect a sudden and sharp increase in the incidence of missing answers to the SAMEMP question, starting in January 2007 followed by a further acceleration through 2009, which never reversed and continued growing gradually through 2015. We identify one important change in survey methodology phased in starting in January 2008 by the




Figure 2: EE^m : Missing answers to the SAMEMP question in eligible (employed in both months) records

US Census Bureau, the Respondent Identification Policy (RIP), which directly impacts the validity of the answer to the SAMEMP question. In a nutshell, the RIP gives, for privacy reasons, the respondent the option not to share their answers, including their employer name that makes the SAMEMP question possible, with any other household members who might happen to answer the survey in subsequent months. A significant number of respondents exercise that option, automatically generating a missing answer to the SAMEMP question a month later. We provide evidence of a very strong selection on unobservable characteristics that correlate positively with EE mobility. We also detect another source of measurement error, similar to but of different nature than the RIP, affecting all CPS cohorts in 2007, and possibly phased out as the RIP was introduced in 2008 and early 2009. This may be related to RIP pre-testing. For all these reasons, observed EE transitions after 2007 poorly estimate the true incidence of EE reallocation.

Based on this evidence, we propose a selection model and a set of identification assumptions, on which we build a procedure to impute missing answers to the SAMEMP question, thus EE transitions, both before and especially after January 2007. Implementing our procedure, we estimate an aggregate EE time series, which differs substantially, over the last 14 years, from Fallick and Fleischman’s (2004), plotted as a dark (blue) line in Figure 1.²

²We make available at <https://campuspress.yale.edu/moscarini/data/>, and will regularly update, the EE time series that we estimate based on both Fallick and Fleischman’s and our methodology, and that

Specifically, our series resets the cyclical peak to early 2008, in line with evidence from administrative quarterly data reviewed later, and reduces the cyclical drop by about half, with a full recovery by 2016, followed by a mild decline thereafter. Thus, our imputed series restores a closer congruence between EE transitions and the business cycle, greatly reduces their cyclical volatility, and eliminates the appearance of a “quit-less recovery” after the Great Recession and of declining EE dynamism in the US labor market since the early 2000s. We also present the first empirical evidence of the large negative impact of the COVID-19 crisis on the pace of EE reallocation; our imputed EE probability series drops even more dramatically than the Fallick and Fleischman (2004) series, as predicted by the selection model, because response rates to the SAMEMP questions rose sharply, reflecting the observed higher availability of previous survey respondents under home lockdown.

The paper is organized as follows. In Section 2 we illustrate the features of the monthly CPS designed to detect individual EE transitions, with a detailed description of the pertinent SAMEMP question. In Section 3 we present our new empirical evidence of the sudden increase in the incidence of missing answers to this question starting in 2007, and relate it to the introduction of the Respondent Identification Policy by the Census Bureau around that time. In Section 4 we provide evidence that the RIP significantly changed measured EE transitions. In Section 5 we propose and implement an imputation procedure of missing answers, hence of EE transitions, based on a model of selection by unobservable worker characteristics that affect the propensity both to answer the survey and to change the job. In Section 6 we compare our imputation results with those from other datasets, we address the impact of survey attrition, and we examine the impact of the COVID-19 pandemic. Brief conclusions take stock of the results and highlight open issues in the measurement of labor market transitions in the CPS, that we leave for future research.

2 EE transitions: data and baseline measurement

2.1 The Current Population Survey (CPS)

The CPS is a monthly survey of about 60,000 households, which has been conducted by the Census Bureau for the Bureau of Labor Statistics for more than 60 years. The information that allows us to detect employer changes has been available only since the 1994 survey

we plot, smoothed, in Figure 1, as well as the time series based on a Missing at Random assumption. Figure 14(a) plots all three times series, not smoothed.

redesign, as described below.³

Despite not being primarily intended for longitudinal analysis, the CPS contains a panel component and can be used to follow individuals over short periods of time. In each month the full CPS sample is divided into eight “Rotation Groups,” with each housing unit being interviewed for four consecutive months, then removed from the sample for an eight-month period, and finally interviewed for another four months. Hence, in any month, one-eighth of the sample households are interviewed for the first month (i.e., the first Rotation Group), one-eighth are interviewed for the second month, one-eighth for the third month, etc.

The CPS has several advantages and disadvantages over panel datasets, such as the Panel Study of Income Dynamics and the National Longitudinal Survey of Youth, in studying labor market states (employment/unemployment, occupation, industry) and related transitions. The first advantage is the large number of individuals in the sample. The second advantage is the high frequency of observations over time, as the CPS is conducted monthly, as opposed to panels that conduct yearly interviews about the entire history of the previous 12 months. The monthly frequency minimizes (although does not eliminate) time aggregation problems due to multiple within-period undetected transitions and to the respondent’s incorrect recall of past events. The third advantage is the wealth of information about demographics, which compares well with that of proper panel data. Finally, only the monthly CPS is updated in a timely manner every month, which makes it uniquely useful to policymakers.

Since the CPS samples housing units (i.e., addresses) and not families or individuals, attrition can occur for one of three main reasons: temporary absence (hospitalization, imprisonment, vacation), migration (to go to college, to enlist in the military, to form a family, to follow or to separate from a spouse, and for work-related reasons, including retirement), and mortality. Thus, the main disadvantage of the CPS is that some attrition is potentially correlated with EE transitions. In Section 6, we provide evidence that this correlation is in practice very weak: most people move for non-job related reasons. In contrast, panel datasets track individuals wherever they move, although they too suffer from significant attrition because of their longer time span. The Survey of Income and Program Participation shares many desirable features of the monthly CPS, but the lower interview frequency (every four months until 2014 and yearly since then) generates recall error in reports and significant delay in the release of new data. Another disadvantage of the CPS is the very limited longitudinal dimension, as individuals are followed for eight (non-consecutive) months, as opposed to decades for panel surveys. This is an unavoidable consequence of the much richer

³Most of the overview information presented in this section is directly based on the official description of the CPS at the Bureau of Labor Statistics website (<https://www.census.gov/programs-surveys/cps/data/datasets.html>).

information set provided by the CPS: since so many questions are asked again every month, they can be asked only for a short period of time, lest becoming harassment.

2.2 Matching monthly CPS files

Matching monthly CPS files means uniquely identifying records in consecutive survey months that refer to the same individual. In principle, the re-interviewing process in the monthly CPS should allow us to match three-fourths of the sample in any given month to the next month, while one-fourth of the sample exits due to rotation (though individuals in their fourth month can be linked eight months forward). As mentioned, however, various kinds of attrition reduce the fraction of individuals that can actually be matched.⁴

The relevant question to identify the transition of interest in this paper, from employer to employer, was introduced as part of the CPS redesign in January 1994. Therefore, we focus on post-1993 data. For the period through April 1995, our matching procedure follows the traditional methodology that combines ID variables (numerical identifiers assigned by the Census Bureau) with some observable individual characteristics, such as age, gender, and race, because there are multiple identical IDs within the same monthly file. As is well known in the literature, between May 1995 and August 1995 matching is impossible due to unavailable ID variables. Thus our analysis cannot cover those four months. Starting in September 1995, the Census Bureau ensures that ID variables are unique, making it unnecessary, and even harmful, to use observable characteristics in establishing the unique matches. In general, matching probabilities are fairly high, although over the past several years attrition grew by about two percentage points. Details are in Appendix A.1.

2.3 The 1994 survey redesign: Dependent Interviewing

An overhaul of the interviewing technique took place in 1994.⁵ Before then, *every* month, respondents were asked anew: (i) for whom they worked, (ii) what kind of business that was, (iii) what kind of work they were doing, (iv) what their most important activities were, and (v) what sector they were working in. This information was later used by CPS staff to assign employer, occupation and industry codes to each individual. This “Independent Coding” procedure had at least two serious shortcomings. First, asking these questions was very cumbersome for the interviewer, and respondents typically complained about answering the same questions repeatedly. Second, and more important for our purposes, asking these

⁴Madrian and Lefgren (2000) and Feng (2001) evaluate in depth the design of the matching criteria of annual (March) CPS records. They build on earlier work in Welch (1993) and Peracchi and Welch (1995).

⁵This description is based on Polivka and Rothgeb (1993). See also Moscarini and Thomsson (2007).

questions independently every month introduced a significant amount of spurious shifts in occupation and industry. Indeed, in a small validation study of occupational coding based on company records and employees' descriptions of their own tasks, Mathiowetz (1992) finds that CPS staff coded occupations incorrectly about half the time when not told that two consecutive records concern the same individual. More remarkably, when told that the two records did come from the same individual, these expert coders still found a 12% disagreement rate between the company record and the employee's description of their tasks.

To reduce the interview burden and misclassification, in 1994, the Census Bureau introduced a number of changes to the survey. The most important change for our purposes is "Dependent Interviewing" (which implies "Dependent Coding"). For those individuals who are reported being employed both last and this month, the interviewer asks the following additional question regarding their main job, that we referred to as "SAMEMP":

- IODP1

Last month, it was reported that (name/you) worked for (company name). (Do/Does) (you/he/she) still work for (company name)?

- Yes
- No

If the answer is No, then this is followed by questions about occupation with the new employer, which is then coded independently of the previous one. If the answer is Yes, then two more questions follow, asking to confirm the description of activities given a month before. If everything is confirmed, then Dependent Coding applies and automatically assigns the same occupational code as in the previous month.

As a result, it has become standard to start the time series of the average EE transition probability in 1994, exploiting answers to the SAMEMP Dependent Interviewing question. We will follow this approach. Note that the SAMEMP question is retrospective and only asked of individuals who are employed in both the past and current month. Therefore, in order to compute the share who answer No, and estimate the average EE probability, in principle we do not need to match records, but can just use cross-sections. In practice, the dataset reports a missing answer to SAMEMP for one of three reasons: the individual was not employed a month earlier; the respondent declined to answer; or, despite the individual being employed in both the past and current month, the record was not eligible for Dependent Interviewing, as explained later. Disentangling these reasons is crucial and only possible by matching records, because the dataset does not provide this information directly.

One last set of individuals remains out of reach: those who were employed in the past

month and, in the meantime, changed address and left the survey in the current month, thus cannot be matched. As explained earlier, taking another job is one of the many possible reasons for moving home. When the entire household moves out, another household often moves into the same address, possibly for the same reason as the outgoing one, including household members taking a new job. But any labor market transition will be missed both for outgoing and incoming households at the time of the move. This is an unavoidable limitation of an address-based survey, which will always lead to underestimating the average EE transition probability. In Section 6, however, we provide abundant empirical evidence that this bias is quantitatively negligible.

3 Missing answers to the SAMEMP question

3.1 Facts

Within the matched records between month t and $t - 1$, those that are employed in both months are eligible for the SAMEMP question in month t . Throughout the paper, whenever we mention “eligibility,” we refer to this criterion and, unless otherwise explicitly stated, analyze this eligible sample. In this sample, we count those who answer No to this question. The ratio between this count and the total number of employed in the initial month within the matched sample is our measure of the EE probability.⁶

The highest hurdle in this apparently straightforward computation is caused by missing answers to the SAMEMP question among eligible records. Those missing answers cannot contribute to the numerator of the EE probability: although we know that these people are employed in both months, we do not know whether at the same company or not. The question is whether the true, unobserved answer was positive or negative. The issue is real even for small percentages of missing answers, because the raw monthly EE probability, computed by just discarding records with missing answers, is small (around 2%), and we do not know the conditional EE probability among those missing answers. For example, suppose that only 1% of all answers are missing but that they are all EE movers in truth. Then, the true EE probability would increase by one half, from 2% to roughly 3%.

In Figure 3, the higher, darker (blue) line illustrates how the share of eligible records with missing answers to the SAMEMP question (EE^m) evolved since the introduction of Dependent Interviewing. Four facts stand out. First, this share has always been positive and non-negligible. Second, it has been rising over time. Both facts were already noticed

⁶Note that the denominator includes some individuals who are no longer employed in the current month.





Figure 3: EE^m : Missing answers to the SAMEMP question by rotation groups in eligible (employed in both months) records

by Moscarini and Thomsson (2007), Figure 3, who at the time analyzed data through 2006. Third, we see a dramatic and persistent jump in January 2007. This fact is new, and cause for great concern. There are further visible sharp accelerations through early 2009. Fourth, the share drops visibly in April and May 2020 and then rebounds, following the COVID-19 shock.


Next, the other two, lower lines split the overall EE^m sample into the first four (1-4) and the second four (5-8) Rotation Groups, normalizing by the same total number of respondents who are eligible for the SAMEMP question, thus the two series add up to the higher blue line. Both series show jumps in January 2007. The former (RG1-4) also jumps at the beginning of 2008, and the latter (RG5-8) at the beginning of 2009, explaining the sharp accelerations in the aggregate measure.

Fallick and Fleischman (2004) pioneered the use of Dependent Interviewing to calculate this EE probability, and their time series has become the main reference in the profession. We reconstruct their time series, using their described methodology and assuming, as they do, that missing answers to the SAMEMP question are stayers.⁷ Our reverse-engineered time series and the one that Fallick and Fleischman make available on their websites coincide

⁷This assumption is not described in Fallick and Fleischman (2004), but was confirmed in a private communication with Charles Fleischman, whom we thank.



(a) Fallick and Fleischman (2004) replication



(b) vs. Missing at Random

Figure 4: EE probability: Fallick and Fleischman (2004) Series

Notes: Due to the missing observations between May 1995 and August 1995 in the raw series, the 12-month trailing moving averages are available only after September 1996.

almost perfectly, as we show in Figure 4a, where the two lines lie on top of each other.⁸ Their treatment of missing answers is potentially problematic even before 2007, more so since then. Instead of treating missing answers as stayers, one can assume that the EE probability of these missing answers is the same as that among valid responses. In Figure 4b, we can see that this Missing-At-Random (MAR) assumption brings the level of the EE probability up noticeably. The gap has been widening since around 2007, in line with increasing incidence of non-response shown in Figure 3.

The natural question is: what happened in 2007-2009? We now provide evidence that the likely culprit is another seemingly small change in the CPS interview protocol, the Respondent Identification Policy (RIP), introduced around that time.

3.2 The Respondent Identification Policy (RIP)

Polivka et al. (2009) provide the following description: “The Respondent Identification Policy (RIP) is the Census Bureau policy that prohibits the sharing of information with other household members unless the person who originally provides the information consents to

⁸Fallick and Fleischman (2004) also exclude Rotation Groups 1 and 5 from their calculations, to avoid the so-called “first rotation group bias,” and focus on transitions between months in sample 2-3, 3-4, 6-7, and 7-8. We follow them to replicate their series in Figure 4. In the rest of our analysis, however, we include all Rotation Groups, including 1 and 5, thus transitions between months in sample 1-2 and 5-6, because we find that they make little difference to the aggregate time series, but they increase the sample size for our imputation procedure of missing answers to the SAMEMP question, described later.

the sharing.” They also describe the cognitive testing that was performed before rolling out the RIP, in order to find the phrasing of the relevant question that would be correctly understood by the maximum number of respondents. The final formulation:

- We will recontact this household next month to update this information. If we are unable to reach you and we talk to someone else instead, is it OK if we refer to the information you gave us?
 - IF NEEDED: An example of this type of question is: “Last month (name) was reported as a teacher. Is (s/he) still a teacher?”
 - IF NEEDED: It will help make the next interview go faster

was still misunderstood by a significant minority of tested respondents.

The CPS Interviewing Manual (April 2015) describes the RIP in Chapter 2.D: “If the original respondent, which we refer to as the ‘RIP respondent,’ wishes their information to be confidential, and they are not available for a subsequent interview, you cannot conduct dependent interviewing. However, if the RIP respondent permits you to verify their information with anyone in the household, then you can conduct dependent interviewing. [...] The instrument will only allow one person to be the RIP respondent. Once the RIP question is asked and the RIP respondent is selected, the RIP question will not be re-asked in subsequent months. You may change the answer to the RIPFLG question during the initial interview only. The only time the RIPFLG will change in subsequent interviews is when there is a replacement household.” Therefore, once the RIP is implemented, a negative answer by the first RIP respondent invalidates dependent interviewing for the entire 4+4 month sequence of that household, unless the household moves out of the address and is replaced by another one moving in. For this reason, although “Any household member 15 years of age or older is technically eligible to act as a respondent,” the Manual then continues: “If at all possible, try to interview the most knowledgeable member of the household. In most situations, this individual will be the reference person or the spouse of the reference person.” In turn, the reference person is defined as “The first person mentioned by the respondent, who either owns or rents the ‘sample unit’ (e.g., house, apartment).”

Polivka et al. (2009) also report that the RIP question is not asked in single-person households, while 14.4% of the RIP questions that were asked for all of 2008 received a negative answer, from respondents who are observationally different from the population. One concern for our purposes is that employed and job-mobile respondents are more likely to answer No to the RIP question, suggesting that they have some confidentiality concerns about their work situation, primarily about their earnings. Polivka et al. (2009) also report that,

in 2008, following one of the 14.4% negative answers to the RIP questions, the respondent changed in only one in nine (11%) households in the following month’s interview. Multiplying the two shares, the No response to the RIP question should result in a share of invalid dependent interviewing of just about 1.5%. We showed much larger numbers than this, *especially* after 2008, because a No answer to the initial RIP question has ramifications that propagate to all other household members and beyond the month of the answer and the one following it, and suppresses information.

From now on, our strategy proceeds in three steps. First, in the remainder of this section, we estimate the timing and mode of introduction of RIP in monthly interviews. The variable RIPFLG flags when an interview is subject to the RIP, and contains the answer to the RIP question for that household, but is not available in the public use data, nor in any confidential version of the data that we are aware of. To determine when and how the RIP was rolled out, we thus proceed indirectly. Based on Polivka et al. (2009), our prior is that the RIP was introduced in 2008. To validate this prior, we exploit the fact that the RIP invalidates some answers to the SAMEMP question. Then, we measure the occurrence and size of month-over-month changes in the share of missing answers to the SAMEMP question, EE^m , starting in 2006. We do this for each cohort and rotation group. Because the RIP applies if the household has more than one member and the household member who answers from the second month on differs from the original RIP respondent, we dig deeper into the pattern of EE^m missing answers to the SAMEMP question among eligible records, breaking it down by single-person household and by respondent status (Self/Proxy). Consistent with our assumption, respondent groups that are expected to be more affected by the RIP show the largest jumps in EE^m . We identify the calendar months of these jumps.

Our second step, in Section 4, exploits the exogenous variation across groups in the timing of the RIP introduction, to identify whether the RIP, or something else, caused a change in measured EE transitions. It is highly unlikely that other changes, especially in the labor market, affected those rotation groups exactly in those months and in that same order. Thus, we use a “treatment-control” approach to document that, every time the RIP was rolled out for a group of respondents, i.e., the share of valid answers to the SAMEMP question suddenly declined, so did the measured EE probability among the remaining valid answers, only for that specific rotation group. So changes in the incidence of EE^m cause simultaneous drastic changes in measured EE, which is the object of interest.

In the third and final step, having demonstrated the causal effect of the RIP on measured EE, we attempt to offset it by imputing EE mobility to eligible records with invalid answers to the SAMEMP question, both pre- and post-RIP periods.

3.3 Identification of survey respondents

The CPS is a monthly, addressed-based, household survey. A household is the collection of individuals who co-habit in the same dwelling, i.e., who live and eat together. Every month, a household member answers the survey for all members, including her/himself. Therefore, a specific answer to a question concerning a specific individual can have one of two respondent statuses: Self (S) if the question concerns the respondent and Proxy (P) if it concerns someone else in the household. Over two consecutive months, the respondent may change, and information about a given individual present in the household and in the survey in both months can follow one of five possible sequences of respondent status: SS, SP, PS, PP, and finally PP'. The last sequence indicates that both responses about this individual were given by different Proxies. PS, SP and PP are only possible in households who have at least two members, and PP' at least three members. Because the RIP is triggered by respondent status, and change thereof, we need to identify these sequences.

For this purpose, we use the indicator variable (PUSLPRX) that indicates whether the person answered the survey that month for the household, to identify the respondent (PULINENO) for each household (HRHHID and HRHHID2). We then construct a flag taking values SS, SP, PS, PP, and PP', and we assign it to the "second" observation in each month's EE sequence. That is, the answer to the SAMEMP question in month t is flagged, say, PS if that answer was given by a Proxy in month $t - 1$ and by the individual him/her-Self in month t . Single-person households are easily identified and necessarily belong in the SS group. Figure 5 plots the shares of the five groups in the population of eligible (employed both last and this month) matched records in each calendar month. The shares of SP and PS are virtually identical. We can see that the share of each group is roughly constant until around 2007, and then SS and PP start rising, presumably reflecting the Census Bureau's effort to secure the same respondent in consecutive interviews after the roll-out of the RIP. All shares exhibit sharp temporary blips in 2020.

In principle, the RIP is more likely to affect SP, PS and PP' records, when the identity of the respondent changes from the last month to the current one and is more likely to differ (surely differs in the second month interview) from the identity of the RIP respondent in the first rotation. In this case, should a respondent deny permission to share his/her answers with future, different respondents, Dependent Interviewing after a change of respondent is ruled out, and the answers to several questions, including the SAMEMP question, will automatically be missing. In Figure 5, the sum of SP, PS and PP' estimates the respondent turnover rate. In 2008 this is about 20%, significantly higher than the 14.4% reported by Polivka et al. (2009) after the RIP question was asked to the first Rotation Group. Figure




Figure 5: Shares of previously employed by respondent status over two consecutive months




Figure 6: Missing answers to the SAMEMP question, by respondent status

6 plots the shares of EE^m by respondent status, namely, the proportion of eligible records *within each respondent group* that has no valid answer to the SAMEMP question. These shares rise over time in each group. Consistently with the logic of the RIP, since 2007, these shares are lower (more valid answers) when the respondent's identity does not change (SS, PP) and higher when it changes and the person in question responds neither time (PP').




Figure 7: Missing answers to the SAMEMP question among Self responses

Because the RIP is relevant only when the identity of the respondent within the household changes, we expect an increase in EE^m more when Proxies are involved than for SS records. The upper and darker (blue) line in Figure 7 is identical to the darker (blue) line in Figure 3; unlike Figure 6, we now normalized the number of missing answers by that of all eligible records, in *all* respondent groups. The middle (orange) line in Figure 7 plots the EE^m incidence among the SS group. Even among these SS respondents, there is a small but noticeable jump in EE^m in 2007. This jump, however, largely disappears, when we condition on Self responses throughout all available interviews (the lighter, yellow line), rather than just a pair of adjacent months. When the RIP respondent is P, a negative answer to the RIP question invalidates later SS records. Consider sequences PSSS with EE^m in the second interview. In this sample, during 2010-2016, the incidence of EE^m in the third and fourth interviews, which are classified as SS (resp., $PSSS$ and $PSSS$), is enormous, over 70%.

The recent COVID-19 crisis offers additional evidence in support of our hypothesis that the RIP affected the non-response rate to the SAMEMP question. In Figure 5, we can see that respondent turnover drops drastically in April and May 2020: the shares of SS and PP rise by about eight percentage points. Presumably, people were suddenly more available to respond again to the CPS because they were forced by the lockdown to stay at home. Accordingly, in Figures 6 and 7, the shares of missing answers to the SAMEMP question fall suddenly and dramatically, returning in two months to levels not seen in about a decade.




Figure 8: Missing answers to the SAMEMP question in single-person households (exempt from the RIP) around the time of RIP roll-out.

3.4 Timing of RIP roll-out

Let $\text{RIP}_{i,t} \in \{0, 1\}$ indicate whether the RIP applies to the survey respondent who answers questions regarding individual i in month t ,⁹ and $\text{DI}_{i,t} \in \{0, 1\}$ a valid answer to the SAMEMP (Dependent Interviewing, retrospective) question regarding individual i in month t . Note that i refers to the identity of the person who is the subject of the questions, not to the identity of the respondent. If $\text{RIP}_{i,t} = 1$, so the survey asks the RIP question and the answer is No, then the SAMEMP question cannot be asked and $\text{DI}_{j,s} = 0$ for all members $j \neq i$ of the same household and all calendar months, including t , when the household is interviewed. But it is also possible that the SAMEMP question can be asked and yet the respondent refuses to answer, or does not know the answer, in which case too we have $\text{DI}_{i,t} = 0$.

Let $\Pr(\text{DI}_{i,t} = 0)$ denote the probability of an invalid answer to the SAMEMP question among eligible records in month t , which can be estimated by the observed share of invalid answers EE_t^m . Note that $\text{DI}_{i,t}$ is an individual-level variable, while EE_t^m is an aggregate time series, a population share, whose time series is plotted in Figure 3. Let $\Pr(\text{RIP}_{i,t} = 1)$ be the probability of a record in month t being subject to the RIP. While we do not observe $\text{RIP}_{i,t}$, we are extremely confident that $\Pr(\text{RIP}_{i,t} = 1) = 0$ before 2007 and $\Pr(\text{RIP}_{i,t} = 1) = 1$

⁹That is, $\text{RIP}_{i,t}=1$ whenever individual i at time t is part of a household whose first rotation respondent (not necessarily i) was asked the RIP question and gave an answer stored in the Census variable RIPFLG.

starting sometime in 2009, given the evidence in Figures 1b and 3 and the description in Polivka et al. (2009). Then we estimate before 2007

$$\Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t} = 0) = \Pr(\text{DI}_{i,t} = 0) = \text{EE}_t^m$$

and after 2009

$$\Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t} = 1) = \Pr(\text{DI}_{i,t} = 0) = \text{EE}_t^m.$$

To estimate the object of interest, $\Pr(\text{RIP}_{i,t} = 1)$, in the intermediate period, we use the identity:

$$\Pr(\text{DI}_{i,t} = 0) = \Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t} = 1) \Pr(\text{RIP}_{i,t} = 1) + \Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t} = 0) \Pr(\text{RIP}_{i,t} = 0)$$

and make the following identification assumption: $\Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t})$ is constant over time for either $\text{RIP}_{i,t} = 0$ or 1 in a period of time surrounding the RIP roll-out, 2006-2010, so we can estimate $\Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t} = 0)$ for t in the roll-out period 2007-2009 with the average of $\Pr(\text{DI}_{i,t} = 0) = \text{EE}_t^m$ in months $t \in 2006$ and $\Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t} = 1)$ with the average of $\Pr(\text{DI}_{i,t} = 0) = \text{EE}_t^m$ in months $t \in 2010$. Then, using our estimate $\Pr(\text{DI}_{i,t} = 0) = \text{EE}_t^m$, the last equation can be solved to obtain an estimate of the incidence of the RIP in every month t in the roll-out period between January 2007 and December 2009:

$$\Pr(\text{RIP}_{i,t} = 1) = \frac{\text{EE}_t^m - \frac{\sum_{\tau \in 2006} \text{EE}_\tau^m}{12}}{\frac{\sum_{\tau \in 2010} \text{EE}_\tau^m}{12} - \frac{\sum_{\tau \in 2006} \text{EE}_\tau^m}{12}}.$$

In words, we assume that the entire increase in the incidence of missing answers to the SAMEMP question in this interim period is due to the introduction of the RIP, and is proportional to the share of records introduced to the RIP. We perform this estimation for each rotation group separately.

To further refine our estimate of the interim period, we zoom into the period surrounding 2004-2013, and add two more pieces of information. First, we examine the time series of $\Pr(\text{DI}_{i,t} = 0)$ for single-member household, who are not subject to the RIP, and thus are never asked that question. Figure 8 shows a jump in January 2007, which reverses in February 2008. Therefore, calendar year 2007 is different. Also, after 2007, there is no trend. This is in contrast to the average population, hence to multi-member households, who are vulnerable to the RIP. Therefore, their rising trend in non-response rate must then be related to the RIP. Second, we break down the time series of missing answers to the SAMEMP question not only by respondent status (SS, PS etc.), as done in Figure 6, but also by rotation



Figure 9: Missing answers to the SAMEMP question by Respondent Group: SS (top) and PS (bottom), and by starting Rotation Group: 1-3 (left) and 5-7 (right).

group. To save on space, in Figure 9 we only present results for the SS and PS group, as PP and SP are (resp.) similar. The SS group (top row) shows modest upward jumps in the incidence of missing answers in January 2007, possibly reversed in early 2008, like single-person households. Rotation Groups 5-7 also show a jump in early 2009, staggered in order of rotation (RG5 jumps first, then RG6 a month later, etc.). Conversely, the PS group (bottom row), more likely than SS and single-person households to be affected by the RIP, shows small jumps in January 2007 and huge jumps, again upwards, in January 2008 for RG1-3 and 2009 for RG5-7, again staggered in order of rotation.

We conclude that the RIP was introduced in a staggered manner, by rotation group, starting in January 2008, while during the entire 2007 calendar year some other change in interviewing procedure affected *all* records. We can only speculate on its nature, possibly

a testing phase of the RIP, but we cannot treat it in the same way as we do the RIP, because its impact is clearly different, as it affects households whose characteristics (such as single-person, or SS respondents) make the RIP irrelevant.

Table 1 provides an overview of our estimated timing of the RIP roll-out period. The date in each cell represents the survey start month (cohort) and the first column gives the calendar time. All cohorts and rotation groups are subject to some unknown factor that causes a temporary increase in the incidence of missing Dependent Interviewing answers in the calendar year 2007, indicated by the light shaded area. The RIP is introduced by CPS cohort, starting with the one that entered the survey in January 2008. All new cohorts from that point on are exposed to the RIP (darker shaded area). The RIP roll-out is completed in April 2009, when the last cohort not exposed to it (December 2007) exits the survey.¹⁰

Figure 10 plots the same EE^m share series as in Figure 3, but with respect to the *cohort* dates, i.e. the dates when each cohort entered the survey that are the entries in the table, rather than with respect to the *calendar* dates. We can clearly see a large jump in the January 2008 cohort, as well as a jump in late 2005 followed by gradual increases toward January 2007. This pattern is consistent with Table 1. The oldest cohort that is exposed to the unknown source of measurement error in January 2007, in their last month in sample, is the October-2005 cohort (right upper corner of Table 1), so only one-eighth of that cohort was subject to that error (only in their last rotation). The November-2005 cohort had two interviews subject to that error; the December-2005 cohort had three interviews; the January-2007 cohort had all eight interviews, and this remains the case for all cohorts until December-2007. So when we plot EE^m by cohort (as in Figure 10), EE^m rises only gradually from October 2005 through January 2007. After that, it remains roughly constant during 2007, until the January-2008 cohort, when the RIP is introduced to that cohort and subsequent ones for all eight rotations, with a much more dramatic impact on EE^m .

4 Impact of the RIP on measured Employer-to-Employer transitions

The RIP has the potential to affect measurement of many variables of interest in the monthly CPS. In this paper, we focus on its impact on EE transitions through the non-random decline of valid answers to the SAMEMP question and provide evidence that the RIP introduced

¹⁰Our imputation procedure of EE mobility after 2007 will provide additional evidence of the pattern illustrated in Table 1. The “bias” introduced by the RIP in measured EE, that we estimate for each respondent group and that we aim to correct, settles into a perfectly regular seasonal pattern after 2008, while it is more erratic in 2007. See Figure A.3

Table 1: The RIP introduction pattern

Calendar date	Rotation Group							
	1	2	3	4	5	6	7	8
2006-1	2006-1	2005-12	2005-11	2005-10	2005-1	2005-12	2004-11	2004-10
2006-2	2006-2	2006-1	2005-12	2005-11	2005-2	2005-1	2004-12	2004-11
2006-3	2006-3	2006-2	2006-1	2005-12	2005-3	2005-2	2005-1	2004-12
2006-4	2006-4	2006-3	2006-2	2006-1	2005-4	2005-3	2005-2	2005-1
2006-5	2006-5	2006-4	2006-3	2006-2	2005-5	2005-4	2005-3	2005-2
2006-6	2006-6	2006-5	2006-4	2006-3	2005-6	2005-5	2005-4	2005-3
2006-7	2006-7	2006-6	2006-5	2006-4	2005-7	2005-6	2005-5	2005-4
2006-8	2006-8	2006-7	2006-6	2006-5	2005-8	2005-7	2005-6	2005-5
2006-9	2006-9	2006-8	2006-7	2006-6	2005-9	2005-8	2005-7	2005-6
2006-10	2006-10	2006-9	2006-8	2006-7	2005-10	2005-9	2005-8	2005-7
2006-11	2006-11	2006-10	2006-9	2006-8	2005-11	2005-10	2005-9	2005-8
2006-12	2006-12	2006-11	2006-10	2006-9	2005-12	2005-11	2005-10	2005-9
2007-1	2007-1	2006-12	2006-11	2006-10	2006-1	2005-12	2005-11	2005-10
2007-2	2007-2	2007-1	2006-12	2006-11	2006-2	2006-1	2005-12	2005-11
2007-3	2007-3	2007-2	2007-1	2006-12	2006-3	2006-2	2006-1	2005-12
2007-4	2007-4	2007-3	2007-2	2007-1	2006-4	2006-3	2006-2	2006-1
2007-5	2007-5	2007-4	2007-3	2007-2	2006-5	2006-4	2006-3	2006-2
2007-6	2007-6	2007-5	2007-4	2007-3	2006-6	2006-5	2006-4	2006-3
2007-7	2007-7	2007-6	2007-5	2007-4	2006-7	2006-6	2006-5	2006-4
2007-8	2007-8	2007-7	2007-6	2007-5	2006-8	2006-7	2006-6	2006-5
2007-9	2007-9	2007-8	2007-7	2007-6	2006-9	2006-8	2006-7	2006-6
2007-10	2007-10	2007-9	2007-8	2007-7	2006-10	2006-9	2006-8	2006-7
2007-11	2007-11	2007-10	2007-9	2007-8	2006-11	2006-10	2006-9	2006-8
2007-12	2007-12	2007-11	2007-10	2007-9	2006-12	2006-11	2006-10	2006-9
2008-1	2008-1	2007-12	2007-11	2007-10	2007-1	2006-12	2006-11	2006-10
2008-2	2008-2	2008-1	2007-12	2007-11	2007-2	2007-1	2006-12	2006-11
2008-3	2008-3	2008-2	2008-1	2007-12	2007-3	2007-2	2007-1	2006-12
2008-4	2008-4	2008-3	2008-2	2008-1	2007-4	2007-3	2007-2	2007-1
2008-5	2008-5	2008-4	2008-3	2008-2	2007-5	2007-4	2007-3	2007-2
2008-6	2008-6	2008-5	2008-4	2008-3	2007-6	2007-5	2007-4	2007-3
2008-7	2008-7	2008-6	2008-5	2008-4	2007-7	2007-6	2007-5	2007-4
2008-8	2008-8	2008-7	2008-6	2008-5	2007-8	2007-7	2007-6	2007-5
2008-9	2008-9	2008-8	2008-7	2008-6	2007-9	2007-8	2007-7	2007-6
2008-10	2008-10	2008-9	2008-8	2008-7	2007-10	2007-9	2007-8	2007-7
2008-11	2008-11	2008-10	2008-9	2008-8	2007-11	2007-10	2007-9	2007-8
2008-12	2008-12	2008-11	2008-10	2008-9	2007-12	2007-11	2007-10	2007-9
2009-1	2009-1	2008-12	2008-11	2008-10	2008-1	2007-12	2007-11	2007-10
2009-2	2009-2	2009-1	2008-12	2008-11	2008-2	2008-1	2007-12	2007-11
2009-3	2009-3	2009-2	2009-1	2008-12	2008-3	2008-2	2008-1	2007-12
2009-4	2009-4	2009-3	2009-2	2009-1	2008-4	2008-3	2008-2	2008-1

Note: The date within each cell indicates the survey start month (cohort date). Lighter shades indicate that survey respondents in the cohort are subject to an unknown source of measurement error in Dependent Interviewing. Darker shades indicate that respondents in the cohort are subject to the RIP at that date.




Figure 10: Missing answers to the SAMEMP question by CPS cohort (month of entry into the CPS)

a strong selection. Figure 11 plots the average EE probability of each respondent group, computed under the MAR assumption (i.e. under the assumption that the EE probability is independent of whether or not there is a valid answer to the SAMEMP question). EE probabilities differ very significantly across groups, so the changing composition by group of valid answers to the SAMEMP question, documented earlier, affects in itself the aggregate EE probability. More importantly, now PS and SP are no longer equivalent. The former has a much higher EE probability than SP, which is instead similar to PP. Note that these two-month respondent groups only seldom include the initial RIP respondent, so PP may be affected by the RIP if the first respondent was S (or a different proxy P’).

To further corroborate our claim that the RIP affects measured EE transitions, we run “treatment/control” and “placebo” experiments, which quantify the jumps that are visually manifested in Figure 11 and demonstrate that these jumps are not, for example, due to unusual seasonality. Specifically, using only the sample of valid SAMEMP answers from January 2006 to March 2009, we regress the individual EE dummy on dummies for calendar month and rotation group, and on two treatment dummies, which mark the two shaded areas in Table 1. These two treatment dummies are interacted with respondent status dummies (SS, PP, PS, SP, PP’). The first treatment dummy equals one if an observation is in the light-shaded area of Table 1, which flags the measurement problem of unknown source, and zero otherwise; and the second treatment dummy equals one if an observation is in the




Figure 11: EE transition rates by respondent status (12-month MA)

dark shaded area of Table 1, which indicates exposure to the RIP, and zero otherwise. The observations in 2006 are in the control group and subject to neither of the measurement problems. This regression estimates, for each respondent status group, and controlling for seasonality and the rotation group, the differences in average EE probabilities of the two treatment groups relative to that of the control group. For the “placebo” experiment, we take the 2005-2006 sample and “treat” 2006 observations with a RIP placebo. That is, the (placebo) RIP dummy takes zero for the 2005 sample and one for the 2006 sample. Basically, we estimate in a flexible manner a correlation between absence of answers to the SAMEMP question and negative answers among the valid ones.

Panels (a) and (b) of Figure 12 summarize the impacts of the two dummies. Both of those dummies are associated with significantly lower EE probabilities for all respondent types except SS. For the first treatment, the largest impact is observed among PP, which one can also notice in Figure 11. The RIP treatment results in further declines in EE probabilities. Interestingly, EE probabilities among SS are little affected by either treatment. Our placebo regression (Panel (c)) shows no indication that similar declines are observed a year earlier.



Note that these are not strictly speaking treatment/control regressions, because the experiments are not simultaneous and therefore the effects of the treatments can be confounded with other time effects. In particular, the treatment periods include the Great Recession, which officially started in December 2007. This is particularly problematic for the second dummy (i.e., the RIP dummy), which marks observations in 2008 and early 2009 (see dark-

shaded area in Table 1). For a genuine treatment/control regression we focus on the period between January 2008 and March 2009. One (randomly selected) half of this sample is subject to the first measurement problem and the rest are subject to the RIP. The treatment dummy equals one for those exposed to the RIP and zero otherwise. For the first calendar month of the sample (January 2008), only the January 2008 cohort (the first Rotation Group) is subject to the RIP, and remaining cohorts are not. For the last calendar month in the sample (March 2009), all rotation groups except the last one (December 2007 cohort) are subject to the RIP. This sample structure allows us to identify the effect of the RIP (in addition to the month effect and the rotation group effect) relative to the 2007 unknown measurement issue, controlling for the time effect. Panel (d) of Figure 12 presents the estimated coefficients on the RIP dummy. We can see that the RIP tends to be associated with lower (measured) EE probabilities, particularly among SP and PS, although the effect on the PP' group is not statistically significant.

5 Imputation of Employer-to-Employer transitions

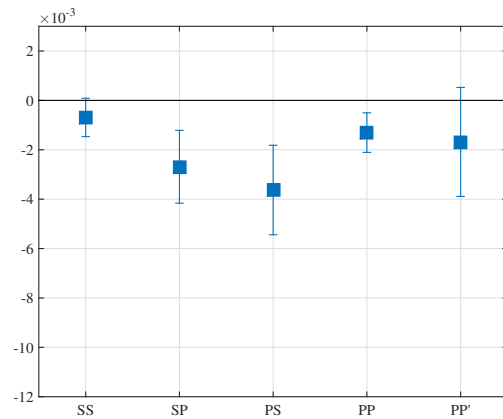

We have just provided empirical evidence that changes in interviewing protocols of unknown nature in 2007 and the introduction of the RIP since the 2008 cohorts altered measurement of the EE probability, differently by respondent status. To redress measurement, we propose an imputation procedure based on data from 1995-2006. A simple approach is to impute EE assuming no selection by unobservable worker characteristics. This will be approximately correct only if observable worker characteristics strongly correlate with the unobservable ones that determine both true EE mobility and the valid answer to the SAMEMP question. Besides demographics, we do have rich observables that arguably do correlate with this type of unobserved heterogeneity, specifically the rotation group, as more job-mobile individuals may be more likely to attrite from the survey and thus no longer answer the SAMEMP question, and the two-month respondent status sequence (SS, PP, PS, SP, PP'), as more job-mobile individuals/households may be more likely to trigger a change in respondent status (SP, PS, PP') and thus the application of the RIP, which prevents the interviewer from asking the SAMEMP question. We will also exploit an aggregate indicator of the labor market prospects for each individual to capture common factors, both trend and business cycle, that affect everybody's true EE probability, independently of the RIP.

If sizable unobserved heterogeneity remains after conditioning on observables, the resulting imputation will not correct for the entire bias in the raw series. Therefore, we introduce a model of selection on unobservables. The difference in average EE probabilities between



(a) Treatment 1 (light-shaded area in Table 1)

(b) Treatment 2 (RIP, dark-shaded area in Table 1)



(c) Placebo (2006 vs 2005)

(d) Treatment 2 (RIP) - Treatment 1

Note: y axis is the change in measured EE probability for that respondent group.

Figure 12: “Treatment and control” regressions (90% CI).

pre- and post-2007 data, given the same observables (worker characteristics, rotation group, respondent status, aggregate indicator), measures the sample selection of those who do answer the SAMEMP question after 2007, when the share of missing answers to the SAMEMP question suddenly rises, and after 2008, when the RIP is implemented. So, for those who do not answer the SAMEMP question, namely for the missing records that we want to impute, the bias is the opposite of this difference, scaled by proportions of valid and invalid records. For example, if individuals who are more affected by the RIP tend to have a *higher* true EE probability, then their selection out of the sample will make the bias in the post-RIP observed EE probability negative, more so the larger the relative incidence of missing records. We now formalize this insight.

5.1 Imputation: model

In order to clarify the possible sources of bias that the changes in the CPS interviewing protocol, especially the RIP, introduced in measuring EE flows, and to obtain a precise imputation formula, we lay out a statistical model. For simplicity, we refer to the RIP as the only source of measurement error after 2007. When we implement this procedure, we treat observations from light-shaded cohorts in Table 1 separately, given our earlier evidence. The same model described below applies to those intermediate observations, after replacing “RIP” with “unknown source of missing answers to the SAMEMP question in 2007-2009,” and the $RIP_{i,t}$ indicator with a dummy for the light-shaded area in Table 1.

Let $E_{i,t}$ denote an indicator function that individual i is employed in month t , with observable characteristics $Y_{i,t}$ (a vector). Recall that $DI_{i,t} \in \{0, 1\}$ indicates a valid answer to the SAMEMP (Dependent Interviewing, retrospective) question, and let $EE_{i,t} \in \{0, 1\}$ indicate an employer-to-employer move (that the valid answer is No). A statistical model is

$$\begin{aligned}\Pr(DI_{i,t} = 1 \mid E_{i,t-1} = E_{i,t} = 1) &= f_{DI}(Y_{i,t}, \theta_{i,t}) \\ \Pr(EE_{i,t} = 1 \mid E_{i,t-1} = 1) &= f_{EE}(Y_{i,t}, \theta_{i,t}),\end{aligned}$$

where θ is an unobservable individual attribute, whose distribution may depend on observables Y . We impose one main assumption on the model: $f_{EE}(Y, \theta)$ is increasing in θ for every Y . This unobserved heterogeneity θ is thus interpreted as the propensity to change job. We are interested in the average mobility of formerly employed workers for each month t , $\mathbb{E}[EE_{i,t} \mid E_{i,t-1} = 1] = \Pr(EE_{i,t} = 1 \mid E_{i,t-1} = 1)$. Some formerly employed workers do not experience an employer-to-employer transition, $EE_{i,t} = 0$, because they separate from their job into nonemployment, $E_{i,t} = 0$. The main issue that we face is that, for the others who stay employed and are thus eligible for the SAMEMP question, we are interested in their average mobility unconditional on a valid answer, $\mathbb{E}[EE_{i,t} \mid E_{i,t-1} = E_{i,t} = 1] = \Pr(EE_{i,t} = 1 \mid E_{i,t-1} = E_{i,t} = 1)$ for each month t , but we only observe the realization of their $EE_{i,t}$ when there is a valid answer $DI_{i,t} = 1$, namely $\mathbb{E}[EE_{i,t} \mid E_{i,t-1} = E_{i,t} = 1, DI_{i,t} = 1] = \Pr(EE_{i,t} = 1 \mid E_{i,t-1} = E_{i,t} = 1, DI_{i,t} = 1)$. The last two expectations do not coincide due to selection on both observables and unobservables into giving a valid answer $DI_{i,t} = 1$. The unobservable individual attribute $\theta_{i,t}$ is assumed to be time-varying. Its persistence captures fixed unobserved traits of individual i , such as preference for job stability, which also determine the person’s propensity to be home to answer the survey, or to give permission to share that information with future respondents under the RIP. Its time variation captures random events, such as receiving a job offer that brings i out of the house for a job interview on the

survey day and triggers a non-response.

In principle, we could specify the functions f_{DI}, f_{EE} of observables Y nonparametrically, i.e., cluster observables in categorical dummies and express each f as a linear combination of such dummies and their full interactions. The number of parameters in, thus the sample size requirements to estimate, such a model would make this strategy infeasible, so we need to impose some parametric structure.

We partition observables Y into two sets $Y = R \cup X$: a “group” R that will be treated nonparametrically, namely, imputation will be performed for each set of individuals in each group separately; and a vector X that will enter parametrically, through regressions using data within each group R . The variables defining the R partition should likely be correlated with unobserved heterogeneity. In our empirical implementation, we define a group by respondent status (SS,SP,PS,PP,PP’), which triggers application of the RIP, which in turn may invalidate eligible records for reasons possibly related to unobserved heterogeneity $\theta_{i,t}$. But, even before the RIP, the $R = PP'$ group exhibits a higher rate of non-response to the SAMEMP question (Figure 6) as well as a higher observed EE probability conditional on valid responses (Figure 11). Therefore, conditioning on respondent group $R \in (SS,SP,PS,PP,PP')$ is useful also before the RIP, as the shares of these respondent groups in the eligible population change over time.¹¹ Note that, in our specific application, a given individual changes respondent group over time depending on the sequence of respondent status over the last two months. The other observables $X_{i,t}$ are discussed below.

To ease notation, from now on we omit the conditioning on employment in consecutive periods, $E_{i,t-1} = E_{i,t} = 1$, hence eligibility to the SAMEMP question, with the understanding that the analysis focuses on this group. Their mobility can then be combined with that (equal to 0) of former employees who no longer work.

We model the probability of an EE transition using the following linear-in- X specification:

$$\Pr(EE_{i,t} = 1 \mid R_{i,t}, X_{i,t}, \theta_{i,t}) = \mathbb{E}[EE_{i,t} \mid R_{i,t}, X_{i,t}, \theta_{i,t}] = \alpha^{R_{i,t}} + X_{i,t}\beta^{R_{i,t}} + \theta_{i,t} \quad (1)$$

with $\theta \mid R, X \sim G(\cdot \mid R, X)$ capturing group-specific unobserved heterogeneity.

Our goal is to estimate the average EE transition rate in the population. By the L.I.E., we can write it as the average of conditional average EE probabilities over respondent groups

¹¹In principle, rotation group is also likely correlated with the individual’s unobserved propensity to change job, because people who move to a different address to take a new job are no longer present in later rotation groups, the well-known issue of geographical attrition in the CPS. In Section 6, comparing with other datasets, we show evidence that survey attrition is quantitatively a minor concern for EE measurement. Defining a group by both 5 respondent statuses and 6 rotation group pairs (1-2, 2-3, 3-4, 5-6, 6-7, 7-8) would require splitting the sample each month into 30 groups, which runs into sample size constraints.

R and observables X :

$$\mathbb{E}[\mathbb{E}E_{i,t}] = \mathbb{E}_{R,X}[\mathbb{E}[\mathbb{E}E_{i,t} \mid R_{i,t} = R, X_{i,t} = X]] \quad (2)$$

so we focus on estimating the conditional rates, and then take their average in the population.

As mentioned, the main issue is that we only observe EE transitions among eligible records that have a valid answer to the SAMEMP question:

$$\begin{aligned} \mathbb{E}[\mathbb{E}E_{i,t} \mid R_{i,t}, X_{i,t}, \text{DI}_{i,t} = 1] &= \mathbb{E}[\mathbb{E}[\mathbb{E}E_{i,t} \mid R_{i,t}, X_{i,t}, \text{DI}_{i,t} = 1, \theta_{i,t}] \mid R_{i,t}, X_{i,t}, \text{DI}_{i,t} = 1] \\ &= \mathbb{E}[\alpha^{R_{i,t}} + X_{i,t}\beta^{R_{i,t}} + \theta_{i,t} \mid R_{i,t}, X_{i,t}, \text{DI}_{i,t} = 1] \\ &= \alpha^{R_{i,t}} + X_{i,t}\beta^{R_{i,t}} + \mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}, \text{DI}_{i,t} = 1] \end{aligned} \quad (3)$$

but do not observe the remaining part of the sample, who do not answer the question:

$$\mathbb{E}[\mathbb{E}E_{i,t} \mid R_{i,t}, X_{i,t}, \text{DI}_{i,t} = 0] = \alpha^{R_{i,t}} + X_{i,t}\beta^{R_{i,t}} + \mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}, \text{DI}_{i,t} = 0]. \quad (4)$$

Selection and bias may occur because the unobserved individual propensity to change job, $\theta_{i,t}$, may be correlated with determinants of obtaining a valid answer to the SAMEMP Dependent Interviewing question ($\text{DI}_{i,t} = 0, 1$) for the same individual, so that

$$\mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}, \text{DI}_{i,t} = 1] \neq \mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}, \text{DI}_{i,t} = 0].$$

If this were an equality, we could impute missing records based only on observables, $R_{i,t}, X_{i,t}$, i.e., projecting observed $\mathbb{E}E_{i,t}$ from the valid answers on these observables and using the regression results to fit the missing answers. In the Appendix, we present the series based on the observables-only imputation: it is nearly identical to the one based on the MAR assumption. Based on this evidence, which contrasts with the drastic change in the pattern of missing answers that we document, we will proceed assuming that the last inequality holds and that we need to correct for this bias.

For this purpose, we make the following **identifying assumptions** about the unobserved component $\theta_{i,t}$ of individual i 's propensity to select into the sample (have a valid answer to the SAMEMP question) and then switch jobs in month t . Later, we describe the imputation algorithm that these assumptions afford.

Assumption 1: No unconditional selection. $\mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}] = 0$.

Given the assumed linear-in- X structure in observables (1), this amounts to assuming that $\mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}]$ is also linear in X , and as such is absorbed in the group fixed effect $\alpha^{R_{i,t}}$

and in the term $X\beta^{R_{i,t}}$.

Assumption 2: No selection before the RIP. Among respondents to the SAMEMP question who are *not* subject to the RIP, unobserved heterogeneity $\theta_{i,t}$ is orthogonal to the validity of the answer to the SAMEMP question, conditional on respondent group $R_{i,t}$ and observables $X_{i,t}$:

$$\begin{aligned} \mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1, RIP_{i,t} = 0] &= \mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 0, RIP_{i,t} = 0] \\ &= \mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}] = 0. \end{aligned}$$

This is a MAR (Missing at Random) assumption about answers to the SAMEMP question within each group $R_{i,t}$ and given other observables $X_{i,t}$. Therefore, before the introduction of the RIP, missing responses to the SAMEMP question are immune from selection on unobservables.

Assumption 3: Time-invariant selection after the RIP conditional on observables. For records subject to the RIP, mean unobserved heterogeneity amongst valid responses to the SAMEMP question is a time-invariant function $b^R(X)$ of respondent group R and observable characteristics X . For all (i, t) :

$$\mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1, RIP_{i,t} = 1] = b^{R_{i,t}}(X_{i,t}).$$

This assumption implies that, within each respondent group R , a valid answer to the SAMEMP question when the respondent is exposed to the RIP may indicate a systematically higher (or lower) mobility than a valid answer to SAMEMP when *not* exposed to the RIP, but this differential mobility only depends on demographics and aggregate labor market conditions gathered in the vector X , and has no other trend nor other time effects. Because we treat the 2007-2009 source of unknown measurement error and the RIP separately, this assumption applies to either, each with its own time-invariant function. Note that, while we assume a time-invariant bias function, the actual bias can change over time for observationally identical individuals because X can contain observable time effects such as trends and business-cycle indicators.

5.2 Imputation: implementation

Our goal is to impute an average EE transition probability to unobserved records as per Equation (4) based only on observables and on our linear model (1) under Assumptions 1-3. This requires estimating α^R , β^R and $\mathbb{E}[\theta \mid R, X, DI = 0]$ for each R, X .

By Assumption 1, taking expectations across i , for every month t

$$0 = \mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}] = \Pr (DI_{i,t} = 0 \mid R_{i,t}, X_{i,t}) \cdot \mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 0] \\ + \Pr (DI_{i,t} = 1 \mid R_{i,t}, X_{i,t}) \cdot \mathbb{E} [\theta \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1].$$

Rearranging, we obtain the key equation on which we build our imputation strategy:

$$\mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 0] = -\frac{\Pr (DI_{i,t} = 1 \mid R_{i,t}, X_{i,t})}{1 - \Pr (DI_{i,t} = 1 \mid R_{i,t}, X_{i,t})} \cdot \mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1]. \quad (5)$$

The strategy consists of estimating all terms on the r.h.s., to obtain from that equation an estimate of the l.h.s., for each record (i, t) , both pre- and post-RIP period. We can then use those estimates in Equation (4) to impute to each missing record an estimated probability of an employer-to-employer move, $\widehat{EE}_{i,t}$. Our final time series is the monthly average of these imputed transitions and of observed $EE_{i,t}$ transitions.

The average “bias” among observed answers, given respondent group and observables, is $\mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1]$, which can be decomposed as follows:

$$\mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1] \\ = \Pr (RIP_{i,t} = 1 \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1) \cdot \mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1, RIP_{i,t} = 1] \\ + \Pr (RIP_{i,t} = 0 \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1) \cdot \mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1, RIP_{i,t} = 0].$$

Now, Assumption 2 implies that $\mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1, RIP_{i,t} = 0] = 0$, and Assumption 3 that $\mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1, RIP_{i,t} = 1] = b^{R_{i,t}}(X_{i,t})$. Next

$$\Pr (RIP_{i,t} = 1 \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1) = RIP_{i,t}$$

is the indicator function that the RIP applies to that record. Crucially, we can assign this indicator based on that record’s CPS cohort from the darker shaded area in Table 1. Combining these implications of our Assumptions 2 and 3, we obtain the following expression for the bias:

$$\mathbb{E} [\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1] = RIP_{i,t} \cdot b^{R_{i,t}}(X_{i,t}).$$

We can now estimate $b^R(X)$ by regressing within each respondent group R the observed EE of those whom we know are treated by the RIP with probability either 0 or 1 on a constant (for α), X (for β) and the interaction of the RIP dummy with a flexible function of X (for $b(X)$). Specifically, for each group $R \in \{SS, SP, PS, PP, PP'\}$ separately, we proceed

through the following **imputation steps**:

1. Using all records eligible for the SAMEMP question ($E_{i,t-1} = E_{i,t} = 1$), every month t run a separate cross-sectional Probit regression of the validity of the answer to the SAMEMP question ($DI_{i,t}$) on observables $X_{i,t}$. Then calculate the predicted value from this regression for each record, and call it $\hat{P}_{i,t}$, an estimate of $\Pr(DI_{i,t} = 1 \mid R_{i,t}, X_{i,t})$ for that individual.
2. Using all available valid answers to the SAMEMP question ($DI_{i,t} = 1$), run an OLS regression of $EE_{i,t}$ on: a constant, $X_{i,t}$, and the interaction of $RIP_{i,t}$ with a flexible function $b(X_{i,t} \mid \gamma)$ parameterized by a vector γ . The resulting estimated coefficients for group R are, respectively, $\hat{\alpha}^R, \hat{\beta}^R, \hat{\gamma}^R$. For all records, predict $\hat{B}_{i,t} = b(X_{i,t} \mid \hat{\gamma}^{R_{i,t}})$, which estimates the bias of valid answers subject to the RIP (the $b^R(X)$ function introduced in Assumption 3, and there assumed to be time-invariant).
3. For each eligible record with missing answer $DI_{i,t} = 0$, impute

$$\widehat{EE}_{i,t} = \hat{\alpha}^{R_{i,t}} + X_{i,t} \hat{\beta}^{R_{i,t}} - \frac{\hat{P}_{i,t}}{1 - \hat{P}_{i,t}} \cdot RIP_{i,t} \cdot \hat{B}_{i,t}.$$

4. Every month t , take the sum of $EE_{i,t}$ when observed ($DI_{i,t} = 1$) and of $\widehat{EE}_{i,t}$ when imputed ($DI_{i,t} = 0$) across all eligible records, so across all respondent groups R and observables X , and divide it by the number of matched individuals in the same CPS cohort who were employed a month before ($E_{i,t-1} = 1$).¹²

By Equation (2), the last ratio is an unbiased (under our model) estimate of the population average probability of transition from employer to employer. Note that the number of non-eligible records of workers who were formerly employed but no longer are ($E_{i,t-1} = 1, E_{i,t} = 0$) contributes to the denominator ($E_{i,t-1} = 1$), but is excluded from the numerator, because they would not contribute to it anyway, by $EE_{i,t} = 0$. Note that the imputation is done for pre-RIP missing records as well, based only on observables: group fixed effect (coefficient α^R) and other covariates X (coefficients β^R). Post-RIP, we also subtract the predicted bias rescaled by the predicted odds ratio of a valid answer, per Step 3 above.

¹²For the denominator, we restrict attention to records that we can match as described in Section 2.2. The retrospective nature of the SAMEMP question allows us to identify also a few records that we cannot match to the previous month, but that have a valid answer, so the Census could match them and know that they were previously employed. Presumably, our failure of matching based on individual identifiers is due to survey processing errors. These cases are so few that they make no difference to the aggregate EE time series of interest, so we feel safe in ignoring them.

We can now illustrate the intuition behind our strategy. The first EE regression in Step 2 exploits Assumption 2 (pre-RIP records are unbiased because Missing at Random and of no selection on unobservables) to compute the X -dependent bias post-RIP, $\widehat{B}_{i,t} = b(X_{i,t} | \widehat{\gamma}^{R_{i,t}})$. In Step 3, Assumption 3 ensures that the function $b(X | \gamma)$ is time-invariant, so $\widehat{B}_{i,t} = b(X_{i,t} | \widehat{\gamma}^{R_{i,t}})$ for the entire post-RIP period. Finally, the smaller the share of missing answers in the survey population, the larger the adjustment in Equation (5) needed to guarantee that unobserved heterogeneity has zero mean in the population by Assumption 1.

A potential concern is that the effect of the RIP may be time-varying, even conditional on respondent group R and on other observables X , violating Assumption 3. Our evidence suggests that this is indeed the case when comparing 2007 and later years, because the behavior of EE^m differs. In the imputation regression, we supplement the $RIP_{i,t}$ dummy with a dummy for the light-shaded area in Table 1, and allow the function $b(X_{i,t} | \gamma)$, specifically the parameter vector γ , to differ between light- and dark-shaded (RIP) areas. So, in Step 2, the regression is run on a constant, observables, two “measurement error” dummies (a light-shaded area dummy and dark-shaded $RIP_{i,t}$) and the interactions of each dummy with a separate flexible function of observables.


5.3 Imputation regressions: specification and results

In Step 2 we specify the function $b(X_{i,t} | \gamma)$ to be linear in the following observables $X_{i,t}$: an aggregate labor market indicator, to be discussed shortly, and dummies for calendar month, Rotation Group (1-2, 2-3, 3-4, 5-6, 6-7, 7-8, denoted by RG1-3 and RG5-7), gender, education (less than HS, HS, Some College, College, Graduate Degree), marital status (Married, Married with Spouse Absent or Separated, Widowed/Divorced, Single), age (16-20, 21-30, 31-40, 41-50, 51-60, 61-70, 71+), major industry (16 major industries, adjusted for breaks to be consistent over time) and major occupation (13 major occupations, adjusted for breaks to be consistent over time). In the Probit regression of Step 1, we omit industry and occupation dummies, because estimation of the full specification sometimes fails to converge.

The aggregate labor market indicator is meant to capture both low-frequency and business cycle variation in the true monthly EE transition probability of the R group, which are unrelated to measurement issues. By absorbing common time variation, this indicator supports the validity of Assumption 3, which requires the RIP bias to remain constant over time. This assumption grows increasingly problematic as time goes by and the pre-RIP period, on which we base our imputation, recedes in the rearview. It is therefore important to verify that no residual trend and cycle are left in the average estimated bias.

For this purpose, we choose as our aggregate labor market indicator the observed average EE probability in the same calendar month of the SS records that are in the first (and second) rotation (EESSRG1). In order not to restrict the effects of the trend and the cycle in such an indicator to be the same, we extract them from the data and enter them separately in the regression. That is, first we fit a quadratic trend to the monthly series of EESSRG1 over the entire period. Then, we incorporate in the vector of observables $X_{i,t}$ both the fitted quadratic trend of EESSRG1 and the deviation from it in month t , as separate regressors. We choose the specific cyclical indicator EESSRG1, to reveal some time patterns in EE transitions for all groups, because this is a reference group that is immune, by design, to the effects of the RIP (SS) and of survey attrition (RG1) on the response rate to the SAMEMP question. After experimenting with many detrending methods, we choose a quadratic trend, and cyclical deviations thereof, because the imputation regression delivers an estimated bias $\widehat{B}_{i,t}$ that, once averaged within each respondent group R , shows no residual trend or cyclical variation (see Figure A.3 in the Appendix), validating Assumption 3. In this sense, EESSRG1 that is filtered through this quadratic R -specific trend performs better than other aggregate indicators of labor market conditions that are also immune from the RIP, such as the UE transition probability of the same R group. We report the results of the Step-2 imputation regression in three tables in Appendix A.2.

We briefly comment on the regression results. In Tables A.2-A.3, RIPFLAG1 and RIPFLAG2 refer, respectively, to the light- and dark-shaded areas in Table 1. As expected, EE mobility is higher among individuals who are in the first rotation group (thus less selected by job-mobility-related survey attrition), less educated, less attached to spouses, and younger. Mobility is also higher when the average EE mobility of RIP-immune records (EESSRG1) is higher, both in trend and business cycle. RIPFLAG1, referring to the intermediate 2007-2009 period, signals a drastic level shift down in observed EE mobility, while RIPFLAG2 ($RIP_{i,t}$) impacts mostly the interaction terms. These two findings indicate that the measurement issue captured by RIPFLAG1 is harder to interpret, while the RIP captured by RIPFLAG2 has no impact on the baseline group and operates mostly through selection. We indeed find that the interactions of the two RIPFLAGs, especially the second ($RIP_{i,t}$), with observables $X_{i,t}$, especially age, are often sizable and statistically significant. The declining age profile of EE mobility, which still survives after controlling for many other worker and job characteristics, is much less pronounced after 2007, and even more so after the introduction of the RIP. This finding indicates that the RIP caused a selection out of the valid sample of more job-mobile individuals among young workers, who are more mobile to begin with. That is, $f_{DI}(R, X, \theta)$ is submodular in age (which is part of X) and unobserved



(a) Missing at Random, 12-month trailing MA

(b) Imputed, 12-month trailing MA

Figure 13: EE probability by CPS Respondent Group

propensity θ to change employer.


5.4 The imputed Employer-to-Employer probability series

In Figure 13, we report, for each respondent group $R \in \{SS, PS, SP, PP, PP'\}$ on which we perform the imputation separately, the time series for the average monthly EE transition probability since 1995, estimated using the Missing at Random assumption (MAR) and our imputation method. All time series are MA-smoothed to remove high frequency noise.


The imputed series, which by construction start diverging from the raw ones after January 2007, are consistently higher, especially for respondent groups SP and PP. This suggests that respondents who denied permission to share their answers with other household members, thus invalidating Dependent Interviewing questions, including SAMEMP, exhibit observable characteristics that strongly correlate in other records with EE mobility.

In Figure 14, we aggregate these group-specific series and report the main result of our paper, which in part replicates Figure 1a: three time series for the average probability of monthly EE transition in the US since 1995, estimated using the Fallick and Fleischman (2004) method (FF), the Missing at Random assumption (MAR), and our imputation method. In the right panel, all time series are MA-smoothed to remove high frequency noise.

By an unfortunate coincidence, measurement issues, as revealed by the January 2007 jump in missing answers to the SAMEMP question, predate by about a year the onset of the Great Recession. Since the EE transition probability is procyclical, the sharp drop



(a) Raw Series



(b) 12-month trailing MA

Figure 14: EE probability: Fallick and Fleischman (2004) vs. Missing at Random vs. Imputed




Figure 15: Comparison with EE probabilities and UE probability (12-month trailing MA)

observed around 2007 in the “raw” (estimated according to either the FF or MAR method) EE probability is easily attributed to the recession. Our imputation procedure leads us to conclude that most of the drop was spurious. While the imputed EE probability did fall, importantly, it declined later, and by much less than the raw EE series and the UE transition probability, which declined by about half starting in late 2008, following the financial crisis. The FF/MAR raw series and our imputed EE series share a weak recovery in 2010-2014, and

then a clear rebound, which ends in 2016. Thereafter, our EE series returns to the pre-Great Recession level of about 2.5% and then starts to mildly decline, while the raw series remains below 2%, generating the false impression of an ongoing long-run decline in this measure of US labor market dynamism. This is another important implication of the imputation. While all measures of firm, job and worker turnover have been trending down in the US in the last few decades, described in concerned terms as “declining fluidity” in the US labor market and “declining dynamism” in US business formation (e.g., Davis and Haltiwanger (2014), Decker et al. (2016), and Molloy et al. (2016)), at least EE turnover appears to have stabilized in the last 15 years.

Note, however, that EE turnover is not an exception in this regard. In Figure 15, we plot two EE probability series (MAR and our series) along with the transition probability from unemployment to employment (the UE probability). One can see that the UE probability and our EE probability series share very similar trends. In particular, in the post-Great Recession period, both series recovered almost fully to the levels immediately prior to the Great Recession. The MAR series, on the other hand, diverged from the UE probability over the same period and the gap has been consistently widening. Furthermore, one can also see that the (local) pre-Great Recession peak of the MAR series came well before the peak of the UE probability (around 2007 in the figure), and fell steadily over the next 5 years or so. Our series, in contrast, exhibits roughly a symmetric hump during the period surrounding the Great Recession (roughly between 2004 and 2010 in the figure) and the UE probability shares a similar symmetric pattern during the same period, although it displays sharper cyclical responses. Importantly, we do not use the UE probability in our imputation at all and there is no reason to believe that the UE probability is also plagued by the measurement issues that affected the measurement of EE transitions. Thus, Figure 15 provides independent evidence that validates our imputation within the CPS.

6 Comparison with other datasets, and the impact of the COVID-19 shock

To further corroborate the validity of our imputation, we compare the average level and time series variation (trend and business cycles) of our CPS-based measures of the EE transition probability with those drawn from other representative datasets of the US labor market. This comparison also offers an opportunity to examine, for the first time, the impact of the COVID-19 crisis on EE reallocation in the US.

6.1 Average levels and survey attrition

As mentioned earlier, a limitation of the monthly CPS for our purposes is its address-based nature. If an employed individual moves out of a selected housing unit to take another job, possibly a household head with the whole household in tow, the survey will lose track of them and miss the EE transition altogether. The same is true of any previously employed individual/household who moves, shortly after, into the same housing unit, thus into the survey, to take another job without any jobless spell. We do not know their employment status before they enter the survey. That is, we need to worry about the correlation between survey attrition and EE mobility.

The CPS classifies non-interviews into three categories. Type A is when the Census interviewer is able to confirm that the same household is living in the unit, but unable to conduct the survey for a variety of reasons. Type B is when the survey unit (the house) is unoccupied and vacant, whether for rent or sale, or held off the market. Type C is when the unit is permanently ineligible; the typical case is “Demolished.” Finally, a Replacement occurs when one household moves out of the unit but is immediately replaced by a different household. We find that, every month since 1994, between 2% and 3% of the records of employed workers who are not in outgoing rotation groups cannot be matched one month forward, because of Type B non-interviews and Replacement. Therefore, the share of movers (out of the address and the survey) among employed workers, whose subsequent employment status is unknown, is comparable in magnitude with our estimated EE transition probability. This makes the impact of survey attrition on EE mobility potentially dramatic. If employed people moved house only to take another job, the true EE transition probability would roughly double our estimate. We can show, however, that this concern is not borne out by other data.

Our first comparison is with the quarterly Longitudinal Employer-Household Dynamics (LEHD), an administrative, matched employer-employee dataset, which contains quarterly reports on total earnings accruing to each US worker from each employer over the entire calendar quarter (U.S. Census Bureau (2021)). Unlike the monthly CPS, this source does not suffer from missing answers, but EE transitions still require an imputation, because of a time aggregation bias. Specifically, we know when a worker earned income from two different employers A and B in quarter t , but to label this an EE transition from A to B in quarter t we have to rule out the possibility that there was a jobless spell in between, which the dataset does not report. Hyatt et al. (2014) propose and implement a filter based on changes in “main employer,” defined as the main source of earnings over two consecutive quarters, and this is the methodology adopted by the Census to estimate the LEHD Job-to-Job Flows

(J2J) series that we use.¹³

To make our monthly CPS estimate of EE comparable in levels with the quarterly LEHD's, we select CPS individuals who have complete interview histories (from RG1-RG4) and no missing answer to the SAMEMP question at any point in the survey. We focus on cohorts who enter the survey from January 2000 through December 2005, which is the early period covered by the LEHD, and when the RIP does not apply yet to the CPS so we can use raw numbers and not our imputed series. In this set, we identify the number of workers who were employed for all four consecutive months RG1-RG4, and estimate the share who experience at least one EE transition during those three pairs of months (quarter). We obtain 5.53%, which is almost identical to the 5.55% average in the quarterly LEHD over the same period. This congruence is reassuring both about Hyatt et al.'s (2014) time aggregation correction in the LEHD and about the irrelevance of geographical attrition in the CPS for EE measurement.

To further corroborate the last point, we turn to the Survey of Income and Program Participation (SIPP). In principle, unlike the CPS, this representative survey tracks individuals even when they move. In practice, the SIPP also suffers from attrition, but at lower rates than the CPS. We select the 2014 panel, when the survey first asked about the reason for a change of address. The 16 possible reasons¹⁴ mentioned to SIPP respondents well illustrate the variety of job-unrelated reasons for people's moves. Of those who were employed at least part of month t , 1.06% moved within state, 0.21% moved to a different state, and 0.01% moved abroad between months t and $t + 1$. Of the within-state (out of state) movers, 4.54% (resp. 30.68%) say they moved to take up a new job. Overall, about 0.1% of those who were initially employed changed address to take up a new job. This further suggests that the bias in the CPS due to correlated attrition and EE mobility is quantitatively negligible.

To check whether the SIPP itself well represents the fraction of employed workers who change address, we use the American Community Survey, an annual, large, representative cross-sectional sample of the US resident population. The IPUMS Abacus tabulates the share of *currently* employed workers who report having moved house in the last year. In 2014-2018 this share, on average and nearly constant, is 12.2% who moved within state, 2.5% from out of state, and 0.5% from abroad. These annual numbers correspond almost exactly

¹³The J2J rate is currently available in 2000:Q2-2019:Q3 from <https://lehd.ces.census.gov/data/>.

¹⁴1. Change in marital/relationship status; 2. To move into own apartment or house; 3. Other family-related reason; 4. New job or job transfer; 5. To look for work or lost job; 6. To be closer to work or school; 7. Other job-related reason; 8. Wanted to own home, not rent; 9. Wanted a better quality apartment or house; 10. Wanted a better neighborhood; 11. Cheaper housing; 12. Other housing-/neighborhood-related reason; 13. Disaster loss (fire, flood, hurricane, etc.); 14. Eviction/foreclosure; 15. Always lived here (never moved); 16. Other reason (specify).




Figure 16: Quarterly EE probability: CPS, LEHD and JOLTS
 Note: Shaded areas indicate NBER dated recessions.

to the monthly numbers from the SIPP reported above.

6.2 Time-series variation

To gauge the cyclical behavior of our imputed series against alternative data sets, besides the LEHD, we also draw from the monthly Job Openings and Labor Turnover Survey (JOLTS), a rotating survey of about 16,000 establishments. The quit rate in JOLTS is the ratio between the number of employees who quit their establishments over the last month, excluding retirements, which are accounted for separately, and the initial level of employment at those establishments. Because JOLTS considers employers, not workers, it cannot distinguish between quits to other jobs and quits to non-employment. It can, however, accurately distinguish between quits and layoffs, because the employer is liable for experience-rated Unemployment Insurance taxes only in the latter case.

To facilitate comparison of our CPS-based series with the seasonally adjusted quarterly series of the LEHD-based J2J series, we seasonally adjust (Census X13) both our monthly CPS series and the monthly JOLTS series through September 2020, take quarterly averages, and rescale them so that the average levels of our CPS-based, LEHD and JOLTS series match up for the first three years of the sample.

Figure 16 reports the results. The CPS-based FF and MAR series start dropping in early 2007, well before the Great Recession, and never recover to pre-recession levels. In

contrast, our imputed series as well as LEHD and JOLTS all drop in earnest during the Great Recession, especially in 2008:Q3 when the financial crisis begins, and all recover their pre-recession levels by 2016. We conclude that declining dynamism in US labor markets ends in the 20th century, at least as far as mobility between firms is concerned. In 2007, our imputed series remains flat, while the other four series are all declining, although not synchronously, suggesting that 2007 remains a challenge for imputation. The correlation coefficient of each of the three quarterly CPS series in Figure 16 with the LEHD series is 0.6521 (FF), 0.7174 (MAR), and 0.8173 (our imputed series). The highest correlation of our series is close to that shared by all three measures pre-RIP (2000:Q2-2006:Q4): 0.8857 (FF), 0.8888 (MAR), 0.8910 (our imputed series), providing further evidence in favor of our imputation. The correlation coefficient of each of the three quarterly CPS series in Figure 16 with the JOLTS quarterly series is 0.4769 (FF), 0.5509 (MAR), and 0.7442 (our imputed series).

During the Great Recession, the LEHD and JOLTS series drop proportionally a lot more than our imputed series, and indeed than any CPS series. This difference in cyclical response raises the concern that our imputation might be over-correcting the drop due to the RIP. We know, however, that, for the LEHD, the large drop is due, at least to some extent, to time aggregation. The described procedure to eliminate in the LEHD spurious EE transitions, which had a short jobless spell in between, is more likely to succeed when short jobless spells are rare, namely, in the trough of the recession. That is, any remaining time aggregation is likely to bias the average level of the LEHD EE transition rate upwards, but this bias is procyclical, as it clearly emerges during the Great Recession due to its severity. Obviously, some modest time aggregation exists also in the CPS, because the SAMEMP question does not distinguish between direct EE transitions and very short jobless spells that complete within the month. Regarding JOLTS, quits to nonemployment are likely to be procyclical, because they are less risky at times of high employment, and thus amplify the cyclical volatility of the overall quit rate in the figure.

Finally, all current series (CPS and JOLTS), once seasonally adjusted and quarterly averaged, drop sharply during the COVID-19 lockdown and the resulting freeze of the US labor market in the second quarter of 2020, and then rebound. As in the Great Recession, the drop is especially pronounced in JOLTS.

In Figure (17) we return to monthly observations and zoom onto 2019-2020. We plot our imputed series and the MAR series, as well as, for reference, the UE transition probability, all seasonally adjusted. The pronounced drop in the EE probability after March 2020 is followed by a strong recovery, which was complete, and more, by summer. EE transitions




Figure 17: EE and UE transitions during the pandemic: Jan 2019 - Dec 2020

slowed down again in Fall 2020, in line with the U.S. macroeconomic recovery, although the level of EE in December 2020 was not unusually low. This pattern indicates that the pandemic dramatically delayed EE reallocation during the year (the .6% drop in the EE probability in early 2020 amounts to about one million fewer workers who changed employer per month), but did not significantly change its total volume. While the EE probability declined, in contrast, the UE probability experienced a huge temporary surge, most likely due to recalls. In this last recession, EE appears to have been a more meaningful real-time gauge of the U.S. labor market than UE. More generally, this graph illustrates how an accurate and prompt measure of high-frequency EE reallocation can inform policy.

6.3 Comparison with LEHD by demographics

The LEHD has been linked to the annual March CPS through the early 2010s (e.g., Hyatt et al. (2018), Bollinger et al. (2019)), so it can be linked to the basic monthly CPS files too, although we found no such instance in the literature. In principle, by doing so we could check the quality of our imputation on error-free LEHD individual records, at least for 2007 through early 2010s. Due to time aggregation in the LEHD, however, this validation would produce asymmetric results, deleting some of the imputed transitions (when the LEHD shows continuity of employment at the same company) but leaving the rest in limbo: even the very careful procedure used to label transitions in the LEHD cannot fully disentangle EE from

EUE, and the types of workers who seem more likely to be affected by the RIP (say, the young) may also be more prone to unemployment. This asymmetry would likely introduce bias.

Here it is critical to emphasize that, from a conceptual viewpoint, even a single day of unemployment in between jobs, which the LEHD could never detect, makes a huge difference to earnings and productivity dynamics, if the worker did not know about the new job when they separated from the old one. From a practical viewpoint, whether a very short non-employment spell should be counted or not for an EE transition is a matter of interpretation. Some employed workers line up a new job and then take some time off before starting it; others lose a job involuntarily, but are lucky enough to find a new one quickly. No dataset allows us to reliably identify this distinction, which is conceptually critical to determine the productivity and earnings implications of the transition. Therefore, researchers have to make some, necessarily dataset-specific, assumptions, which complicate the comparison between datasets with different frequencies. For this reason, we find it more informative to aggregate the CPS to the same quarterly frequency as the LEHD, and to compare the average EE series, which are our main focus. We did it at the national level, but we can dig a bit deeper, into demographics.

The quarterly EE transition probabilities constructed from the LEHD are available also by demographics and some job characteristics from 2000:Q2 to 2019:Q3. Because the monthly CPS has a small sample size, relative to the LEHD, we cannot disaggregate it too finely by demographics as well as time. For example, if we slice the data by all characteristics available for the LEHD series, the corresponding group-specific EE probabilities in the CPS will be very noisy, even when aggregating months into quarters, because of sample size. As a compromise, we form ten demographic groups, by gender and age (19-24, 25-34, 35-44, 45-54, 55 and up). We then estimate the monthly EE probability for each group, and average it for each quarter. We end up with 40 observations per year. For each of the 13 years in the post-RIP period, 2007-2019 included, we use these 40 observations (30 in 2019) to run a separate linear regression of the quarterly EE probability from the LEHD on that from the CPS, across groups and quarters, and a constant, weighting by group employment shares in the LEHD. As we are primarily interested in across-group covariation between LEHD and CPS, we omit group fixed effects. At any rate, within-group time variation over the four quarters of a single calendar year is swamped by variation across groups.

Figure 18 plots the time series of the estimated intercept (left panel) and slope (right panel) of this sequence of year-by-year linear regressions of LEHD on CPS, either MAR or our imputed series, with 95% confidence intervals. It is clear that, relative to the MAR

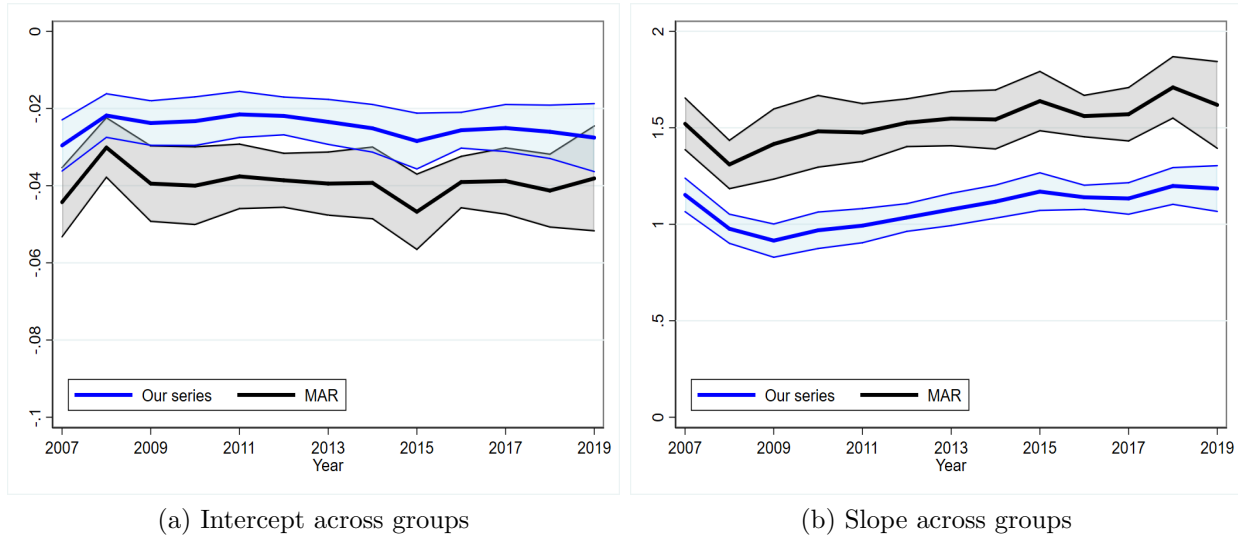


Figure 18: EE probability: relationship between each CPS-based series and LEHD series across demographic groups

series, our imputed series produces an estimated intercept significantly closer to zero and an estimated slope closer to one at all points in time, establishing a closer congruence between our series and the LEHD by demographics. The negative estimated intercept with our imputed series most likely reflects another type of time aggregation: multiple EE transitions within a quarter count as just one in the LEHD, so the level of EE is generally lower than in the CPS. The intercept is even more negative for MAR, due to the downward bias caused by the RIP, which our imputation aims to correct.

7 Conclusions

We measure aggregate employer-to-employer transitions made by workers, without any intervening jobless spell, in US labor markets. We draw from the monthly Current Population Survey. We uncover a drastic increase in the incidence of missing answers to the pertinent survey question (SAMEMP) starting in January 2007, predating by about a year the full introduction of a new interviewing policy, the Respondent Identification Policy (RIP). We provide evidence that these answers are not missing at random, and these interviewing changes caused a serious permanent downward bias in the standard measure of employer-to-employer transitions. We propose a model of selection by observable and unobservable worker characteristics, and build on it to impute the missing answers to recover the true aggregate employer-to-employer monthly transition probability. We show that its decline

observed during the Great Recession started about a year later and was much less dramatic than the raw, biased series indicates, and had fully recovered by 2016, if not earlier. We conclude that the EE transition rate in the US is procyclical, but less volatile and higher than previously thought, and presents no low-frequency trend in the 21st century.

Our analysis still faces important limitations. First, while we are confident that the 2007 observations are corrupted, we are only certain that the RIP was introduced in 2008, so we still do not know the true nature of the 2007 problem, although we apply the same procedure to address it. Second, the share of invalid answers to the SAMEMP question in the CPS was modest but slowly rising even before 2007; at that point, this share experiences a few upward jumps, mostly related to the introduction of the RIP, until early 2009, but then continues to rise after 2009, smoothly but much faster than before 2007. We also show that the share of CPS monthly records that can be matched month-over-month has been declining significantly since 2010 or so. Therefore, underlying trends in response rates have been causing an overall deterioration in the quality of CPS observations, and appear to interact with the RIP. While our imputation procedure addresses some of this trend by controlling for sample composition of the missing Dependent Interviewing answers, it is plausible that additional and progressive selection by unobservables is unfolding, unrelated to the RIP and partially immune to our imputation. In future research, we plan to investigate the causes of these ongoing trends. Getting to the bottom of this measurement issue is especially important in light of the recent debate on declining dynamism in the US labor market.

References

- Bollinger, Christopher R., Barry T. Hirsch, Charles M. Hokayem, and James P. Ziliak**, “Trouble in the Tails? What We Know about Earnings Nonresponse 30 Years after Lillard, Smith, and Welch,” *Journal of Political Economy*, 2019, 127 (5), 2143–2185.
- Burdett, Kenneth and Dale Mortensen**, “Wage Differentials, Employer Size, and Unemployment,” *International Economic Review*, 1998, 39 (2), 257–273.
- Davis, Steven and John Haltiwanger**, “Labor Market Fluidity and Economic Performance,” in “Re-Evaluating Labor Market Dynamics” Jackson Hole Economic Policy Symposium Proceedings, Federal Reserve Bank of Kansas City, 2014, pp. 17–107.
- **and Till Von Wachter**, “Recessions and the Costs of Job Loss,” *Brookings Papers on Economic Activity*, 2011, Fall, 1–72.

- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda**, “Declining Business Dynamism: What We Know and the Way Forward,” *American Economic Review Papers and Proceedings*, 2016, 106 (5), 203–207.
- Fallick, Bruce and Charles Fleischman**, “Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows,” 2004. Federal Reserve Board Finance and Economics Discussion Series 2004-34.
- Feng, Shuaizhang**, “The Longitudinal Matching of Current Population Surveys: A Proposed Algorithm,” *Journal of Economic and Social Measurement*, 2001, 27 (1-2), 71–91.
- Foster, Lucia, John Haltiwanger, and Chad Syverson**, “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?,” *American Economic Review*, 2008, 98 (1), 394–425.
- Fujita, Shigeru and Giuseppe Moscarini**, “Recall and Unemployment,” *American Economic Review*, 2017, 107 (12), 3875–3916.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song**, “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 2014, 122 (3), 621–660.
- Haltiwanger, John, Henry Hyatt, Lisa Kahn, and Erika McEntarfer**, “Cyclical Job Ladders by Firm Size and Firm Wage,” *American Economic Journal: Macroeconomics*, 2018, 10 (2), 52–85.
- Hubmer, Joachim**, “The Job Ladder and Its Implications for Earnings Risk,” *Review of Economic Dynamics*, 2018, 29, 172–194.
- Hyatt, Henry, Erika McEntarfer, Ken Ueda, and Alexandria Zhang**, “Interstate Migration and Employer-to-Employer Transitions in the U.S.: New Evidence from Administrative Records Data,” *Census Bureau Center for Economic Studies Working Paper 16-44R*, 2018.
- , –, **Kevin McKinney, Stephen Tibbets, and Doug Walton**, “Job-to-Job (J2J) Flows: New Labor Market Statistics from Linked Employer-Employee Data,” *JSM Proceedings 2014, Business and Economic Statistics Section*, 2014, pp. 231–245.
- Lentz, Rasmus and Dale Mortensen**, “An Empirical Model of Growth Through Product Innovation,” *Econometrica*, 2008, 76 (6), 1317–1373.

- Madrian, Brigitte and Lars John Lefgren**, “An Approach to Longitudinally Matching Current Population Survey (CPS) Respondents,” *Journal of Economic and Social Measurement*, 2000, *26* (1), 31–62.
- Mathiowetz, Nancy**, “Errors in Reports of Occupation,” *Public Opinion Quarterly*, 1992, *56* (3), 352–355.
- Molloy, Raven, Christopher Smith, Riccardo Trezzi, and Abigail Wozniak**, “Understanding Declining Fluidity in the U.S. Labor Market,” *Brookings Papers on Economic Activity*, 2016, *Spring*, 183–237.
- Moscarini, Giuseppe and Fabien Postel-Vinay**, “On the Job Search and Business Cycles,” 2018. Unpublished Manuscript.
- **and Fabien Postel-Vinay**, “The Job Ladder: Inflation vs. Reallocation,” 2019. Unpublished Manuscript.
- **and Kaj Thomsson**, “Occupational and Job Mobility in the U.S.,” *Scandinavian Journal of Economics*, 2007, *109* (4), 807–836.
- Peracchi, Franco and Finis Welch**, “How Representative Are Matched Cross-Sections? Evidence from the Current Population Survey,” *Journal of Econometrics*, 1995, *68* (1), 153–179.
- Polivka, Anne and Jennifer Rothgeb**, “Redesigning the CPS Questionnaire,” *Monthly Labor Review*, 1993, *116* (9), 10–28.
- **, Polly Phipps, Christine Rho, and Hugette Sun**, “The Current Population Survey’s Experience With the Respondent Identification Policy,” Slides Prepared for May 2009 AAPOR Conference 2009.
- Postel-Vinay, Fabien and Jean-Marc Robin**, “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 2002, *70* (6), 2295–2350.
- Topel, Robert and Michael Ward**, “Job Mobility and the Careers of Young Men,” *Quarterly Journal of Economics*, 1992, *107* (2), 439–479.
- U.S. Census Bureau**, “Job-to-Job Flows (2000-2019) [computer file],” 2021. Washington, DC: U.S. Census Bureau, Longitudinal-Employer Household Dynamics Program [distributor], accessed on December 10, 2020. R2020Q4.

Welch, Finis, “Matching the Current Population Surveys,” *Stata Technical Bulletin*, 1993, 12 (2), 7–11.

Appendix

A.1 Matching CPS files


To match records in January 1994-April 1995, we first take the variable HRHHID, which is 12 digits, and then concatenate it with a 5-digit number, which is in turn created by combining the following three variables: sample number (HRSAMPLE), serial suffix (HRSERSUF) and household number (HUHHNUM). The resulting 17-digit number still does not uniquely identify the household and therefore, even when combined with person line number (PULINENO), the individual. For this reason, following the literature, we also use the individual's age, gender and race to establish an individual match.^{A.1}

Starting in September 1995, HRHHID is 15 digits, and its three additional digits, along with the 5-digit number formed by HRSAMPLE, HRSERSUF, and HUHHNUM as before,^{A.2} generate a 20-digit number that uniquely identifies the household. Individuals within the household can then be identified by PULINENO without using observable characteristics. In fact, after September 1995 these observable individual characteristics are likely to generate “spurious mismatches,” because the Census Bureau occasionally “scrambles” respondents’ age information, and more generally because these characteristics may be measured with error. ID variables are arguably more fundamental to the entire survey and thus mistakes in coding the ID variables are likely to be rare or to be eventually corrected before the data is made public.

Figure A.1a presents the probability that a respondent who appears in the month- t micro data in Rotation Groups 1-3 or 5-7 also appears in the month $t+1$ data. Note that Rotation Groups 4 and 8 in month t are excluded from this calculation, because they rotate out of the survey in the following month as a result of the survey design. The solid line in Figure A.1a gives the matching probability based on ID variables only, while the dashed line gives that based also on the additional three observable characteristics. In general, matching probabilities are fairly high although over the past several years attrition increased by about two percentage points. The difference between the two lines measures unmatched observations due to inconsistencies in either age, sex, or race. One can see that the dashed line exhibits occasional downward spikes (the spike at the end of 1995 is common to both methodologies). In Figure A.1b, we present probabilities that either age, sex, or race is inconsistent between the two months, conditional on IDs matching between the two months.

^{A.1}We allow for age to increase by one year between the two months.

^{A.2}Starting in May 2004, this five-digit part, named HRHHID2, is directly available from the data.



(a) Comparison of Matching Rates

(b) Inconsistencies in Observable Characteristics

Figure A.1: Matching rates

The occasional drops in the dashed line in (a) are mostly due to inconsistencies in the age information, although race also contributed to the drop at the end of 2002, because of changes in the coding of the race variable that occurred between December 2002 and January 2003.

A.2 Imputation Step 2: Regression results

Tables A.1-A.3 summarize the results of the Step-2 regressions in our imputation procedure described on page 30. Table A.1 presents the coefficient estimates for the 1995-2006 pre-RIP sample, Table A.2 for the interactions with the 2007-2009 cohorts affected by the measurement error of unknown origin, and Table A.3 for the interactions with the RIP, in 2008-2020. We comment on these results in the paper.

To validate the key Assumption 3 for our imputation, Figure A.2 shows the fit of a quadratic trend to our aggregate labor market indicator, the observed average EE probability in the same calendar month of the SS records that are in the first (and second) rotation (EESSRG1). The quadratic trend and the deviations from it enter separately the imputation regression. Then Figure A.3 illustrates the average estimated bias for each month and each respondent group. By construction, the bias is zero before 2007. It is clear that 2007 is different from later years, when the average bias settles into a very regular seasonal pattern, with no visible residual trend and cycle, except for a very small decline for the PP group and a slightly hump-shaped pattern for the PP' group. The seasonal pattern of the bias indicates that the underlying seasonal pattern of EE, clearly visible in pre-2007 data, changed

Table A.1: Imputation regression results

	$R = SS$	$R = SP$	$R = PS$	$R = PP$	$R = PP'$
RIPFLAG 1	-0.094*	-0.105	0.079	-0.087	-0.449*
RIPFLAG 2	-0.010	-0.029	-0.038	0.006	0.032
Rotation Gr. 2-3	-0.002***	-0.001	-0.008***	-0.003***	-0.004**
Rotation Gr. 3-4	-0.003***	-0.002**	-0.012***	-0.003***	-0.009***
Rotation Gr. 5-6	0.000	0.003***	-0.001	0.000	0.000
Rotation Gr. 6-7	-0.002***	-0.001	-0.009***	-0.003***	-0.005***
Rotation Gr. 7-8	-0.004***	-0.001	-0.012***	-0.003***	-0.006***
Sex	0.002***	-0.002***	0.001	0.001*	-0.001
Married Spouse Absent	0.004**	0.007***	0.020***	0.010***	0.010***
Widowed/Divorced	0.003***	0.004***	0.014***	0.008***	0.010***
Never Married	0.003***	0.008***	0.021***	0.009***	0.004**
High School	-0.003***	-0.003**	-0.007***	0.001***	0.007***
Some College	-0.001	-0.002	-0.007***	0.005***	0.010***
College	0.000	-0.001	-0.006***	0.003***	0.013***
Graduate	0.001	0.001	-0.006***	0.003***	0.005*
Ages 21-30	-0.025***	-0.023***	-0.016***	-0.016***	-0.015***
Ages 31-40	-0.033***	-0.036***	-0.027***	-0.026***	-0.028***
Ages 41-50	-0.036***	-0.038***	-0.029***	-0.029***	-0.038***
Ages 51-60	-0.038***	-0.037***	-0.032***	-0.026***	-0.041***
Ages 61-70	-0.039***	-0.036***	-0.032***	-0.024***	-0.039***
Ages 71-	-0.040***	-0.033***	-0.033***	-0.021***	-0.039***
EE-SS-RG1 Trend	0.781***	1.106***	1.147***	0.968***	2.229***
EE-SS-RG1 Cycle	0.352***	0.124*	0.152*	0.257***	0.382***
Constant	0.033***	0.032***	0.035***	0.021***	-0.013

Notes: Base groups: Rotation Group 1-2, male, married spouse present, high school dropouts, and ages 16-20. Each regression also includes month dummies, 16 major industry and 13 major occupation dummies in the initial month. The full results are available upon request. The sample period is September 1995 - May 2020. The superscripts *, **, and *** indicate significance at 10%, 5%, and 1%, respectively. t-statistics are in parentheses.

Table A.2: Imputation regression results: RIPFLAG 1 interaction terms

	$R = SS$	$R = SP$	$R = PS$	$R = PP$	$R = PP'$
Rotation Gr. 2-3	-0.001	-0.001	-0.004	0.003**	-0.004
Rotation Gr. 3-4	-0.001	-0.003	0.004	0.001	0.006
Rotation Gr. 5-6	0.000	0.001	0.002	0.002	-0.007
Rotation Gr. 6-7	-0.001	-0.002	-0.001	0.002	-0.006
Rotation Gr. 7-8	-0.002	-0.002	0.001	0.002	-0.001
Sex	-0.001*	0.001	-0.004**	-0.001	0.005
Married Spouse Absent	0.000	-0.002	-0.001	-0.005*	0.001
Widowed/Divorced	0.000	-0.001	0.001	-0.004***	-0.002
Never Married	0.001	-0.007***	-0.008**	-0.005***	-0.006
High School	0.003**	-0.003	0.006	-0.001	-0.010**
Some College	0.003**	0.000	0.006	-0.002*	-0.007
College	0.004***	-0.001	0.006	0.002	-0.008
Graduate	0.003*	-0.005	0.009*	0.001	-0.003
Ages 21-30	0.006	0.003	-0.01	0.001	0.009*
Ages 31-40	0.008*	0.009	-0.002	0.004**	0.002
Ages 41-50	0.011**	0.008	-0.006	0.003	0.010
Ages 51-60	0.010**	0.004	-0.002	0.001	0.008
Ages 61-70	0.013***	0.004	-0.004	-0.002	0.005
Ages 71-	0.010**	-0.001	-0.011	-0.003	0.006
EE-SS-RG1 Trend	4.610*	4.769	-2.798	4.191	23.691*
EE-SS-RG1 Cycle	0.056	0.907	1.375	-0.879**	-1.241

Notes: Base groups: Rotation Group 1-2, male, married spouse present, high school dropouts, and ages 16-20. Each regression also includes month dummies, 16 major industry and 13 major occupation dummies in the initial month. The full results are available upon request. The sample period is September 1995 - May 2020. The superscripts *, **, and *** indicate significance at 10%, 5%, and 1%, respectively. t-statistics are in parentheses.

Table A.3: Imputation regression results: RIPFLAG 2 interaction terms

	$R = SS$	$R = SP$	$R = PS$	$R = PP$	$R = PP'$
Rotation Gr. 2-3	0.000	0.001	-0.002	0.001*	-0.003
Rotation Gr. 3-4	-0.001	0.002	0.001	0.000	0.006**
Rotation Gr. 5-6	0.001**	0.001	0.000	0.001**	0.000
Rotation Gr. 6-7	0.000	0.000	-0.001	0.002**	-0.001
Rotation Gr. 7-8	0.000	0.000	0.000	0.001	0.003
Sex	-0.001	0.000	0.001	-0.001***	0.002
Married Spouse Absent	-0.001	-0.004	-0.002	-0.005***	-0.006
Widowed/Divorced	0.000	-0.002	-0.001	-0.005***	-0.010***
Never Married	0.000	-0.004***	-0.007***	-0.004***	-0.002
High School	0.000	0.003*	0.008***	0.000	-0.003
Some College	-0.001	0.001	0.008***	-0.001*	-0.005*
College	0.000	0.002	0.010***	0.000	-0.006**
Graduate	-0.001	-0.001	0.011***	0.000	0.000
Ages 21-30	0.013***	0.007*	0.004	0.006***	0.008***
Ages 31-40	0.015***	0.013***	0.007*	0.010***	0.012***
Ages 41-50	0.016***	0.013***	0.010**	0.011***	0.018***
Ages 51-60	0.017***	0.009**	0.011**	0.007***	0.019***
Ages 61-70	0.018***	0.006	0.010**	0.005***	0.017***
Ages 71-	0.019***	0.005	0.014**	0.002	0.020**
EE-SS-RG1 Trend	-0.271**	0.348	0.824**	-0.237*	0.238
EE-SS-RG1 Cycle	-0.027	0.038	0.122	-0.029	-0.023
N	5024870	953107	948677	4386096	352191
R^2	0.004	0.007	0.008	0.007	0.001

Notes: Base groups: Rotation Group 1-2, male, married spouse present, high school dropouts, and ages 16-20. Each regression also includes month dummies, 16 major industry and 13 major occupation dummies in the initial month. The full results are available upon request. The sample period is September 1995 - May 2020. The superscripts *, **, and *** indicate significance at 10%, 5%, and 1%, respectively. t-statistics are in parentheses.

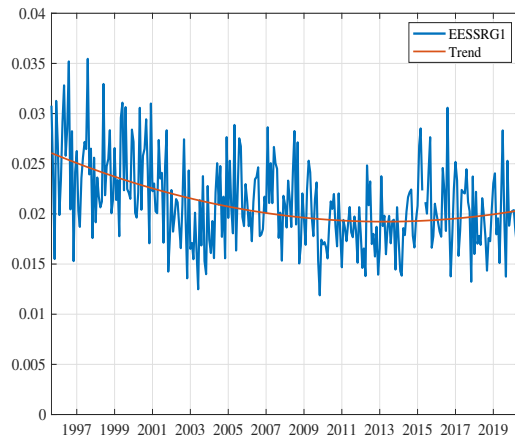


Figure A.2: EE probability and quadratic trend: Self-Self between the first and second months of the survey

permanently with the introduction of the RIP in 2008, and settled into a different, but equally regular, pattern thereafter, further evidence of selection by unobservables. Over the entire post-2007 period, the bias averages to approximately zero for the SS group, and is otherwise negative, reflecting the reduction in measured EE due to the correlation between EE mobility and non-response rate. The bias grows in size moving from the PS to the SP group, then further for PP and is largest for PP’.

A.3 Imputing missing records by observables only

In the main text, we focused on three EE probability series: the Fallick-Fleischman series, the MAR series, and our proposed series. The other obvious possibility is to impute the missing records simply based on observables. That is, we can simply project observed $EE_{i,t}$ from the valid answers on the observables and use the regression results to impute the missing answers. Specifically, we run the imputation regression for each of the five respondent groups (as in our proposed imputation procedure) over three different samples, corresponding to the three shaded areas in Table 1. The latter sample selection is arbitrary, but allowing for the regression coefficients to differ across these three samples appears reasonable. The results are robust with respect to other sample selections as well.






Figure A.3: Estimated average bias $\mathbb{E}_i[\widehat{B}_{i,t}]$

In Figure A.4, we compare this series with the one based on the MAR assumption. The figure clearly shows that imputing the missing records based only on observables results in an aggregate EE probability series that is effectively identical to the MAR series.



(a) Raw Series



(b) 12-Month Trailing Moving Average

Figure A.4: EE probability: Missing at Random vs. Imputed by observables only