Measuring Employer-to-Employer Reallocation*

Shigeru Fujita[†] Philadelphia FED

Giuseppe Moscarini[‡] Yale University and NBER

Fabien Postel-Vinay[§] UCL and IFS

July 2019

Abstract

We revisit measurement of Employer-to-Employer (EE) transitions in the monthly Current Population Survey (CPS). We follow Fallick and Fleischman (2004) and exploit a key survey question introduced with the 1994 CPS redesign. We detect a sudden and sharp increase in the incidence of missing answers to this question starting in January 2007, when the U.S. Census Bureau introduced a relevant change in survey methodology, the Respondent Identification Policy (RIP). We show that the individual records affected by this censoring differ significantly from the rest in terms of observable characteristics, and we provide evidence of selection on unobservable characteristics that correlate with EE mobility. We conclude that observed EE transitions substantially underestimate its true incidence after 2007. We propose a selection model and a procedure to impute missing answers to the key survey question, thus EE transitions, after the introduction of the RIP starting in January 2007. We estimate a different aggregate EE time series than Fallick and Fleischman (2004) over the last 12 years. Specifically, our imputed EE series restores a close congruence with the business cycle, especially with the onset of the Great Recession, exhibits a much less dramatic cyclical drop and a full recovery by 2016, and all but eliminates the spurious appearance of declining EE dynamism in the US labor market after 2000.

*We thank Ann Polivka at the BLS for taking her time with our inquiries about various technical aspects of the CPS, and Henry Hyatt, Jim Spletzer (both at the Census Bureau), and participants in the 2019 Montreal Workshop on Markets with Frictions for comments and discussions. The usual disclaimer applies.

[†]Federal Reserve Bank of Philadelphia, Research Department, Ten Independence Mall Philadelphia, PA 19106-1574, shigeru.fujita@phil.frb.org

[‡]Yale University, Department of Economics, PO Box 208268, New Haven CT 06520-8268, giuseppe.moscarini@yale.edu

[§]University College London, Department of Economics, Drayton House, 30 Gordon Street, London WC1H 0AX, UK, 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 by firms for, employed workers are a natural source of worker bargaining power, and an important determinant of "pure" 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 an individual point of view, that 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 can 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 from nonparticipation (NE) into employment. Since a very significant share of the last two comprises (resp.) recalls by the last employer (Fujita and Moscarini, 2017) and new entry into the labor force, flows that do not directly reallocate employment between firms, EE play a quantitatively dominant role in reallocation of workers across jobs, which is a major driver of aggregate productivity growth (e.g. Foster et al., 2008 and Lentz and Mortensen, 2008). The EE transition probability also appears to be procyclical, but much less volatile than UE probability and unemployment rate. These facts bear significant implications for the cyclical reallocation of labor input between productive activities (firms, industries, 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. We focus on the monthly Current Population Survey (CPS), the official source of the unemployment rate and the premier source of real-time information on labor markets available to policymakers in the US. After its 1994 redesign, the CPS contains an explicit question (variable IODP1) that can be used to identify EE transitions: the interviewer asks whether the name of an individual's current employer is the same as the one recorded in the previous month's interview. In this paper,

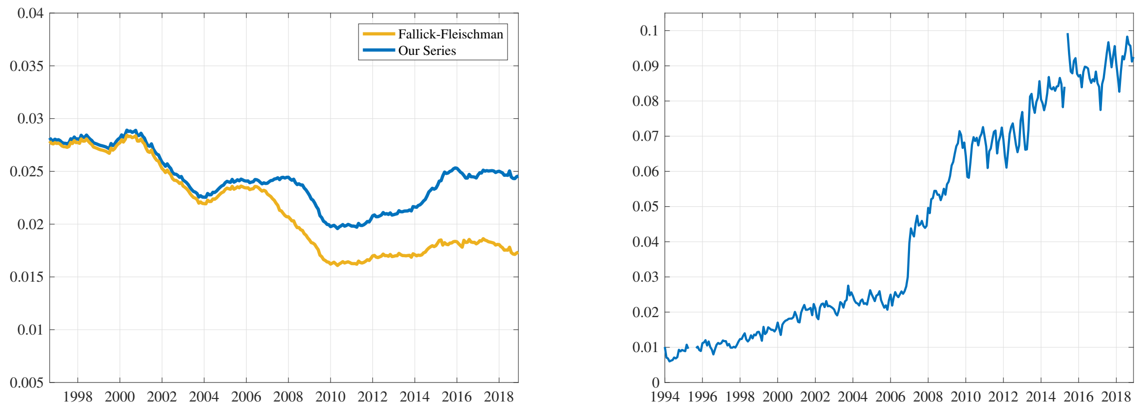

(a) EE transition probability (12-month MA)

(b) Missing answers to question on EE transition

Figure 1: Main Findings

we will refer to this question as "SAMEMP". Fallick and Fleischman (2004) pioneered the use of the answers to this question to measure the average EE transition probability, and provide an estimated monthly time series which 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. We can see a fairly dramatic drop starting in early 2007, which never reverts, thus generating the impression of a strong cyclical drop preceding the Great Recession by a full year, as well as a low frequency downward trend.[1]

In this paper, we revisit measurement of the EE transition probability. Our starting point is Figure 1b. We detect a sudden and sharp increase in the incidence of missing answers to the SAMEMP question, starting in January 2007 followed by a similar one a year later, which never reversed but instead continued growing thereafter. We identify one important change in survey methodology introduced, starting in January 2007, by the 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 her/his answers with any other household members who might happen to answer the survey in subsequent months. A significant number of respondents appear to exercise that option, automatically generating a missing answer to the SAMEMP question a

[1] At this point we should emphasize that the type of transitions we are focusing on in this paper *involve a change of employer* — hence the systematic reference to "Employer-to-Employer (EE) transitions". Those transitions are sometimes referred to as "Job-to-Job" (J2J) in the literature: 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. Those within-employer job changes are excluded from our analysis.

month later. We provide evidence of a very strong selection by unobservable characteristics that correlate positively with EE mobility, therefore concluding that observed EE transitions after 2007 severely underestimate its true incidence.

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 the introduction of the RIP which started in January 2007. Implementing our procedure, we estimate an aggregate EE time series which differs substantially, over the last 12 years, from Fallick and Fleischman's, both plotted in Figure 1a. Specifically, our (imputed) series resets the cyclical peak to early 2008, in line with evidence from quarterly administrative data, and reduces the cyclical drop by about half, with a full recovery by 2016, followed by a flat stretch since then. Thus, our imputed series restores a closer congruence between EE transitions and the business cycle, greatly reduces its 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 after 2000.

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 both the propensity to answer the survey and to change job. Brief conclusions take stock of the results and highlight other 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 Bureau of the Census for the Bureau of Labor Statistics for more than 60 years. The information that allows to detect employer changes is only available since a 1994 survey redesign, as described below.[2]

[2] Most 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 (www.bls.census.gov/cps).

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. Since the interviewers follow 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.

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 completely eliminate) time aggregation problems due to multiple within-period undetected transitions and 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.

The main disadvantage of the CPS for our purposes is its address-based nature, which aligns attrition with geographical mobility, potentially correlated with employer-to-employer mobility. In contrast, panel datasets continue to track the same individuals wherever they are, although panels too suffer from significant attrition because of their much lower interview frequency. The Survey of Income Program Participation has the same desirable features of the monthly CPS and does not suffer from geographical attrition, 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 information set provided by the CPS: since so many questions are asked again every month, they can be asked for a short period of time, lest becoming harassment.

2.2 Matching of monthly CPS files and sample characteristics

Matching monthly CPS files means 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). However, various kinds of attrition reduce the fraction of individuals that can actually be matched. Madrian and Lefgren (2000) and Feng (2001) evaluate in depth the design of the matching criteria of annual (March) CPS records.[3]

Matching individuals across months requires to uniquely identify them. Matching is traditionally achieved by combining ID variables (numerical identifiers assigned by the Census Bureau) with some observable individual characteristics, such as age, gender, and race, because there could be multiple identical IDs within the same monthly file. As we will see shortly, 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 with age, gender, and race. 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. We now provide some details.

To match records in January 1994-April 1995, we first take the variable HRHHID, which is 12 digit, 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.[4]

Starting in September 1995, HRHHID is 15 digit, and its three additional digits, along with the 5-digit number formed by HRSAMPLE, HRSERSUF, and HUHHNUM as before,[5] 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

[3]Both of these articles are inspired by earlier work in Welch (1993) and Peracchi and Welch (1995).
[4]We allow for age to increase by one year between the two months.
[5]Starting in May 2004, this five-digit part, named HRHHID2, is directly available from the data.

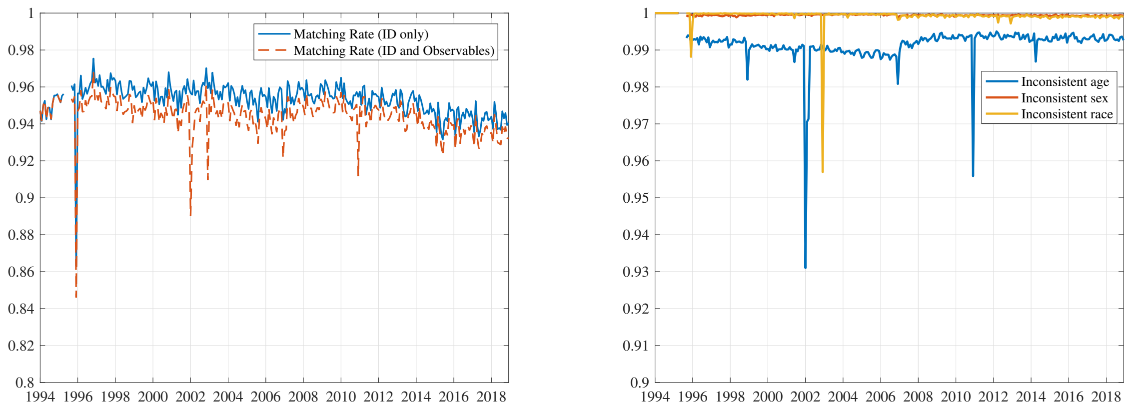

(a) Comparison of Matching Rates

(b) Inconsistencies in Observable Characteristics

Figure 2: Matching Rates

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 2a present 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 2a 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 2b, we present probabilities that either age, sex, or race is inconsistent between the two months, conditional on IDs matching between the two months. 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.

2.3 The 1994 Survey redesign: Dependent Interviewing

An overhaul of the interviewing technique took place in 1994.[6] Before 1994, *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 change of 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. Secondly, and more importantly for our purposes, asking these questions independently from month to 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 a roughly 50% error when occupations are coded without telling the coders that the two records concern the same individual. More remarkably, when told that the two records did come from the same individual, expert coders still found a 12% disagreement rates between the company record and the employee's description of the employee's task, although each coder knows that they should agree.

To reduce the interview burden and the possibility of misclassification, in 1994 the Census Bureau introduced a number of changes in the CPS. The most important change for our purposes is "Dependent Coding" (sometimes referred to as "Dependent Interviewing"). Those that are interviewed in successive months and are reported as employed both last and this month are asked 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 in 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.

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

As a result, it has become standard to start the time series of EE probabilities in 1994, exploiting answers to the SAMEMP Dependent Interviewing question. We will follow this approach. Note that the SAMEMP question is retrospective. Therefore, in order to compute the *number* of individuals who change job over a month, we do not need to match records and can just use cross-sections. Matching is, however, necessary to know how many of the respondents to a specific month's Survey, specifically to questions about employment status and SAMEMP, were employed a month before, thus how many of the previous respondents stayed on the same job, changed job, or stopped working for whatever reason, thus the *probability* of an employer-to-employer transition.

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 question. Unless otherwise explicitly stated, our analysis concerns this eligible sample. Out of 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.[7]

The most formidable 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 this is 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 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 darker (blue) line illustrates how the share of SAMEMP-eligible records with missing answers to the SAMEMP question (EE^m) evolved since the introduction of Dependent Interviewing. Three facts stand out. First, this share has always been present and non-negligible. Second, this share has been rising over time. Both facts were already

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

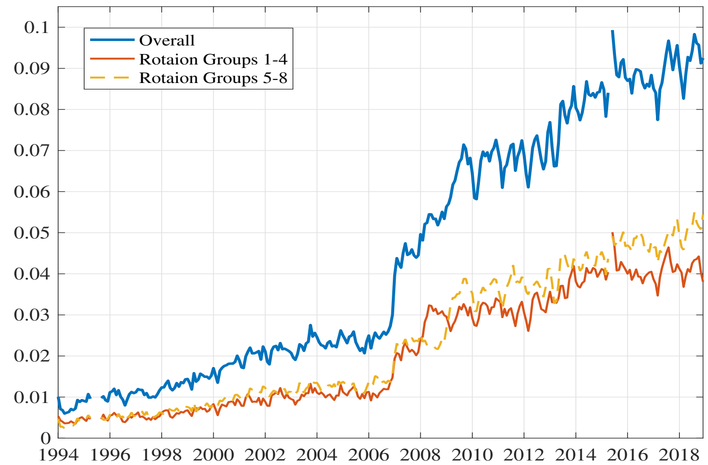

Figure 3: EE^m: Share of missing answers to the SAMEMP question in eligible (employed in both months) records

noticed 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 is also a visible sharp acceleration in both January 2008 and January 2009. Next, the other two lines split the overall EE^m sample into the first four (1-4) rotation groups and the second four (5-8) rotation groups. Both show jumps in January 2007. The former also one in January 2008, and the latter in January 2009, explaining the sharp accelerations in the aggregate measure. Both series are normalized by the same total number of respondents that are eligible to the SAMEMP question (i.e., those that are employed in both months), thus the two series add up to the blue line.

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 blank answers to the SAMEMP question are stayers.[8] Our reverse-engineered time series and the one that Fallick and Fleischman make available on their websites almost perfectly coincide, as we show in Figure 4a, where the two lines lie on top of each other.[9]

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

[9] Fallick and Fleischman (2004) also exclude from their computations the rotation groups 1 and 5, to avoid the so-called "first group rotation 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.

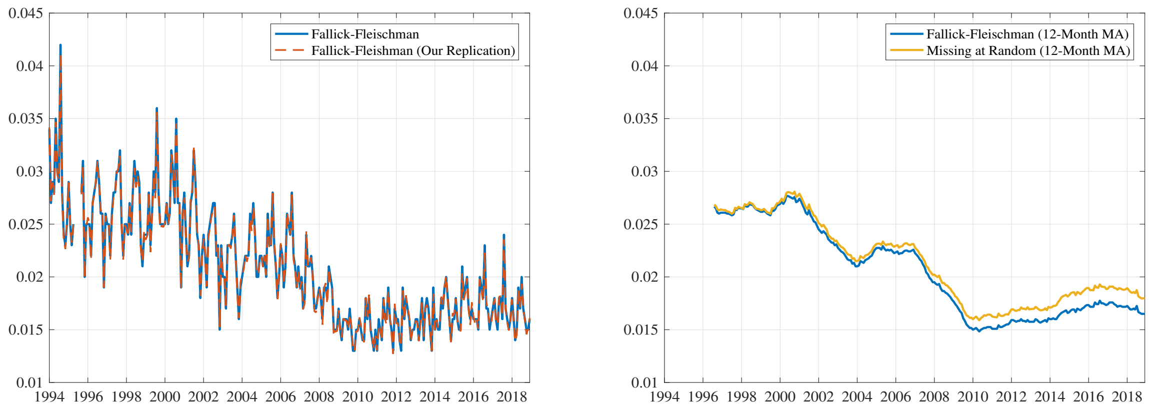

(a) FF Replication

(b) FF vs. Missing at Random (12-month MA)

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.

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 within the valid responses. In Figure 4b, we can see this Missing-At-Random (MAR) assumption brings up the level of the EE probability noticeably. The difference has been expanding since around 2007, as implied by increasing missing records shown in Figure 3.

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

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 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. Polivka et al. (2009) also report that 14.4% of the RIP questions asked in 2007-2008 received a negative answer, and that the sample of such respondents is observationally different from the population. One concern for our purposes is that employed 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. From now on, our strategy proceeds in three steps.

First, in the remainder of this section, we identify the exact timing and mode of introduction of RIP in monthly interviews. The variable RIPFLG flags when an interview is subject to the RIP, but is not available in the public use data. To determine exactly when and how RIP was rolled out, we thus proceed indirectly. We assume that RIP invalidates some answers to the SAMEMP question, which in fact must happen when the previous respondent answers No to the RIP 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, around 2007, the time when we know the RIP was introduced. We do this for each cohort and rotation group. We find two types of jumps in the incidence of EE^m, one roughly twice as large as the other, the former permanent and the latter followed by another jump of similar magnitude. Based on these observations, our working hypothesis is that the RIP was introduced in a staggered manner, first gradually, into a randomly drawn half of a rotation group, then became the standard for all surveyed individuals; in other words, that no other event or interviewing methodology change can explain the drastic jumps in EE^m that we observe in specific months in 2007-2009. Because the RIP applies only if the identity of the household member who answers the survey changes from month to month, we dig deeper into the pattern of EE^m missing answers to the SAMEMP question among eligible records, breaking it down by respondent status (Self/Proxy). Consistently with our assumption, respondent groups that are expected to be more affected by the RIP show the largest jumps in EE^m.

Our second step, in Section 4, exploits the exogenous variation across groups in the timing of the RIP introduction by the Census Bureau, to gain identification power on the question whether the RIP, or something else, causes 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, and 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 the incidence of EE^m causes simultaneous drastic changes in measured EE, which is the object of interest.

Finally, 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 invalid answers, both pre- and post-RIP periods.

3.3 Timing of RIP roll-out

Let $\mathrm{RIP}_{it} \in \{0, 1\}$ indicate whether the RIP applies to the Survey respondent who answers questions regarding individual i in month t, and let $\mathrm{DI}_{it} \in \{0, 1\}$ indicate 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 identify of the respondent. If $\mathrm{RIP}_{it} = 1$ and the answer to the RIP question is No, then the SAMEMP question cannot be asked and $\mathrm{DI}_{it} = 0$. But it is also possible that the question is asked and yet the respondent refuses to answer, or does not know the answer, in which case also $\mathrm{DI}_{it} = 0$.

Let $\Pr(\mathrm{DI}_{it} = 0)$ 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, that we denote by EE^m_t in keeping with our previous notation. Note that DI_{it} is an individual-level variable, while EE^m_t is an aggregate time series, a population share, whose time series is plotted in Figure 3. Let $\Pr(\mathrm{RIP}_{it} = 1)$ be the probability of a record in month t being subject to the RIP. While we do not observe RIP_{it}, we know that $\Pr(\mathrm{RIP}_{it} = 1) = 0$ before 2007 and $\Pr(\mathrm{RIP}_{it} = 1) = 1$ starting sometime in 2009. Then we estimate before 2007

$$\Pr(\mathrm{DI}_{it} = 0 \mid \mathrm{RIP}_{it} = 0) = \Pr(\mathrm{DI}_{it} = 0) = \mathrm{EE}^m_t$$

and after 2009

$$\Pr(\mathrm{DI}_{it} = 0 \mid \mathrm{RIP}_{it} = 1) = \Pr(\mathrm{DI}_{it} = 0) = \mathrm{EE}^m_t$$

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

$$\Pr(\mathrm{DI}_{it} = 0) = \Pr(\mathrm{DI}_{it} = 0|\mathrm{RIP}_{it} = 1)\Pr(\mathrm{RIP}_{it} = 1) + \Pr(\mathrm{DI}_{it} = 0|\mathrm{RIP}_{it} = 0)\Pr(\mathrm{RIP}_{it} = 0)$$

and make the following identification assumption: $\Pr(\mathrm{DI}_{it} = 0|\mathrm{RIP}_{it})$ is constant over time for either $\mathrm{RIP}_{it} = 0$ or 1 in a period of time surrounding the RIP rollout, 2006-2010, so we

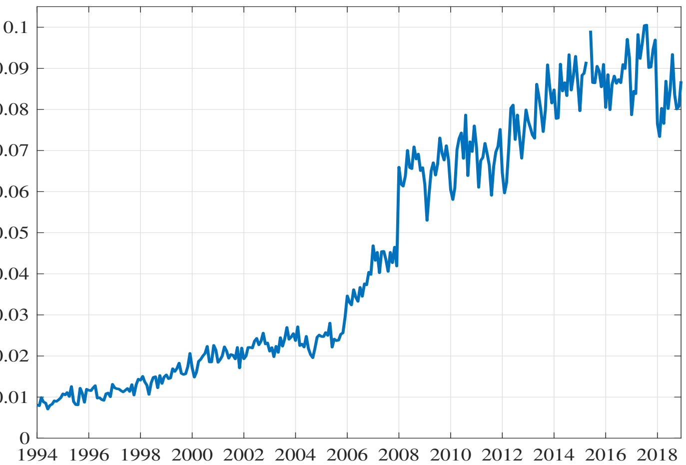

Figure 5: EE^m by Cohort

can estimate $\Pr(\mathrm{DI}_{it} = 0|\mathrm{RIP}_{it} = 0)$ with the average of $\Pr(\mathrm{DI}_{it} = 0) = \mathrm{EE}_t^m$ in 2006 and $\Pr(\mathrm{DI}_{it} = 0|\mathrm{RIP}_{it} = 1)$ with the average of $\Pr(\mathrm{DI}_{it} = 0) = \mathrm{EE}_t^m$ in 2010. Then, using our estimate $\Pr(\mathrm{DI}_{it} = 0) = \mathrm{EE}_t^m$, the last equation can be solved to obtain an estimate of the incidence of the RIP in every month t between January 2007 and December 2009:

$$\Pr(\mathrm{RIP}_{it} = 1) = \frac{\mathrm{EE}_t^m - \frac{\sum_{\tau \in 2006} \mathrm{EE}_\tau^m}{12}}{\frac{\sum_{\tau \in 2010} \mathrm{EE}_\tau^m}{12} - \frac{\sum_{\tau \in 2006} \mathrm{EE}_\tau^m}{12}}$$

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 roll-out of the RIP, and proportional to the share of records introduced to the RIP. We perform this estimation for each rotation group separately.

Table 1 provides an overview of our findings. The RIP is first introduced to one half of the survey respondents in the survey in January 2007, *cross sectionally* to all respondents regardless of rotation groups, so including groups that had already been in the survey for a long time. RIP is then expanded to all respondents, starting in the January-2008 *cohort*. All cohorts from that point on are exposed to the RIP. In the table, the date in each cell represents the survey start month (cohort) and the first column gives the calendar time. The lighter shade indicates that one half of respondents in that cohort are subject to the RIP and the darker shade indicates that all respondents in that cohort are subject to the RIP.

Figure 5 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, rather than with respect to the

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 half of survey respondents in the cohort are subject to the RIP at that date. Darker shades indicate that all respondents in the cohort are subject to the RIP.

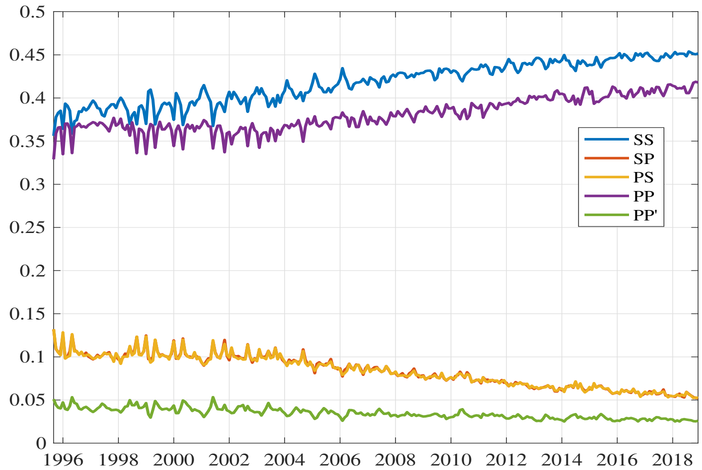

Figure 6: Shares of previously employed by respondent status over two consecutive months

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 RIP in January 2007, in their last month in sample, is the October-2005 cohort (right upper corner of Table 1), so only one half of one eighth of the cohort was subject to the RIP. The November-2005 cohort had two interviews subject to the RIP; the December-2005 cohort had three interviews subject to the RIP, etc. The January-2007 cohort is subject to the RIP for all eight interviews. So when we plot EE^m by cohort (as in Figure 5), EE^m rises only gradually from October 2005 through January 2007. After that there are no jumps, because the January-2007 cohort through the December-2007 cohorts are subject to the RIP for all eight interviews, until January 2008 when the remaining half of the respondents are introduced to the RIP.

3.4 Identification of respondent status

Because the introduction of the RIP appears to be associated with a dramatic change in the share of valid answers to the SAMEMP question, on which EE measurement is based, and because the effects of the RIP depend on the identity of the responder, we need to measure this identity accurately. 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

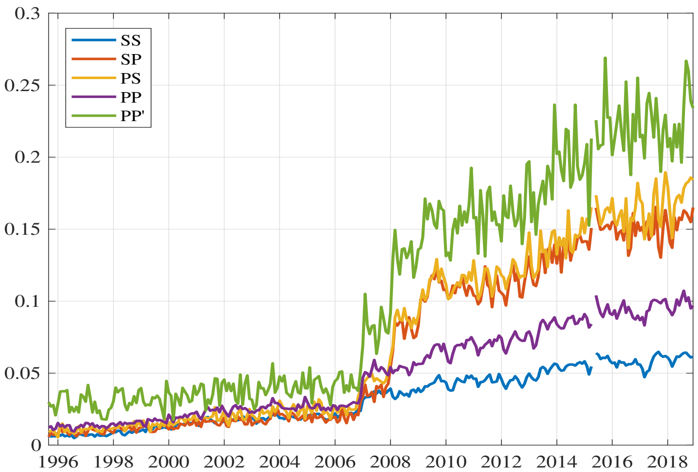

Figure 7: Shares of missing answers to the SAMEMP question, by respondent status

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' which 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) 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 at $t-1$ and by the individual him/her-Self in month t. Figure 6 plots the shares of the five groups in the population of eligible (employed both last and this month) matched records each calendar month. SP and PS are virtually identical. We can see that the incidence of each group is roughly constant until the RIP is introduced, and then the shares of SS and PP start rising, presumably reflecting Census efforts to secure the same respondent as in consecutive interview.

In principle, the RIP can affect only SP, PS and PP' records, when the identity of the respondent changes from last month to the current one. In this case, should a respondent deny permission to share his/her answers with future, different respondents, Dependent

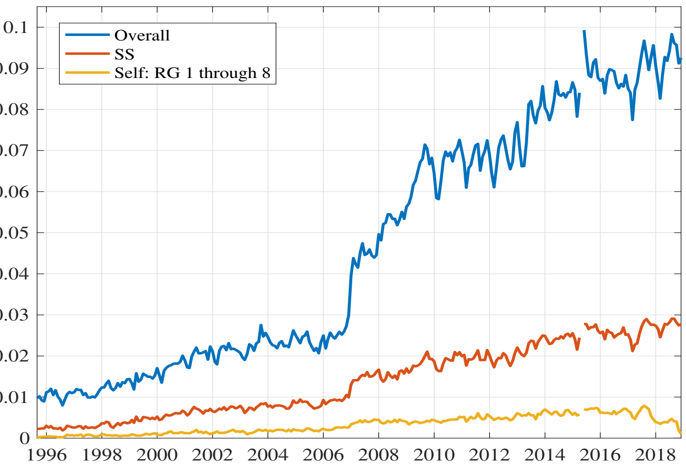

Figure 8: EE^m among Self Responses

Interviewing after a change of respondent is ruled out, and the answers to several questions, including the SAMEMP question, will automatically be missing. Figure 7 plots the shares of EE^m by respondent status. They rise over time in each group. Consistently with the logic of the RIP, these shares are lower (there are 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'). We now discuss reasons why we believe the RIP also affected records of the SS and PP types that should have been immune to it.

Polivka et al. (2009) report that, in 2007-2008 following one of the 14.4% negative answers to the RIP questions, only one in nine respondents need to be replaced. Multiplying the two shares, the No response to RIP should result in a share of invalid dependent interviewing of just about 1.5%. We showed much larger numbers than this, *especially* after the introduction of the RIP, indicating that a No answer to the RIP question has ramifications that propagate beyond the month of the answer and the one following it and suppress information.

Because the RIP is relevant only when the identity of the respondent changes from month to month, we expect an increase in EE^m only when Proxies are involved, not for SS records. The darker (blue) line in Figure 8 is identical to the darker (blue) line in Figure 3. The middle (orange) line plots the EE^m incidence among the SS group, again normalized by the same denominator (matched records employed in both months) as in the earlier figures, so the lowest series in Figure 6, expanding the vertical scale. 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

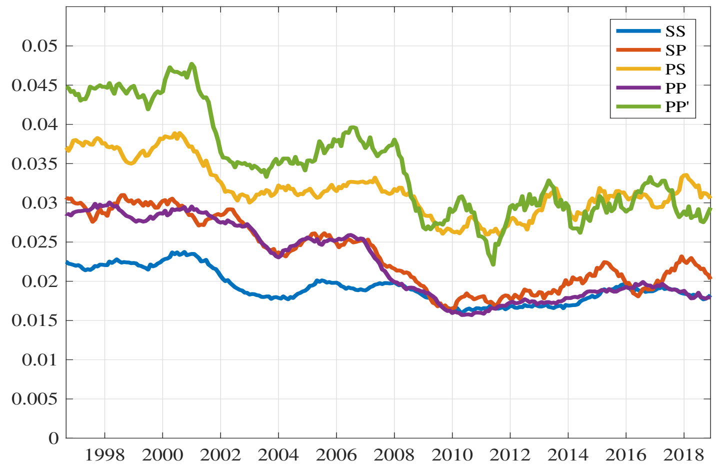

Figure 9: EE transition rates by respondent status (12-month MA)

line), rather than just a pair of adjacent months. Therefore, once a Proxy respondent declines the RIP question in one interview, a small but persistent effect arises on following interviews for the same household, even if answered by Self. That is, the RIP flag makes dependent interviewing difficult also later. Additional evidence is provided by the following observation. Consider sequences P<u>SS</u>S who have EE^m in the second interview (PS in the <u>PS</u>SS sequence). In this sample, in 2010-2016, the incidence of EE^m in the third interview, which is classified as SS, is enormous, over .7, and remains almost as high in the fourth interview, also SS (namely, PS<u>SS</u>), even after P is long gone.

4 Impact of the RIP on measured EE 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 related SAMEMP question, and we provide evidence that the RIP introduced a strong selection. Figure 9 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 the aggregate EE probability just by composition. More importantly, now PS and SP are no longer equivalent. The former has a much higher EE probability than SP, which is instead

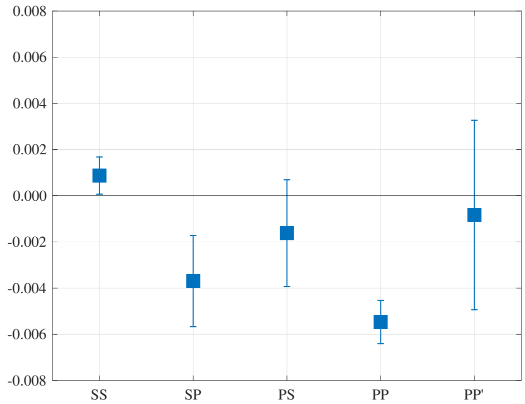

(a) RIP Treatment

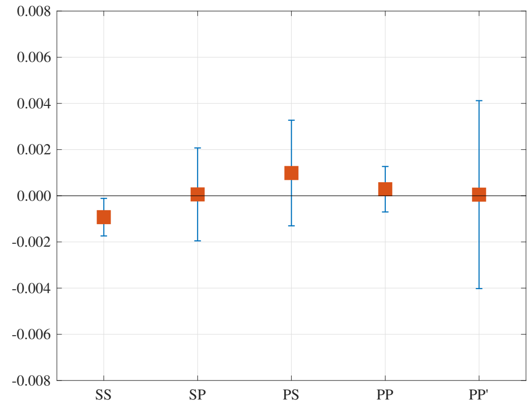

(b) Placebo

Figure 10: Effect of the RIP introduction on measured EE mobility

similar to PP. This suggests that starting with a Proxy answer and/or changing respondent generates more measurement error of employer identity in the first of the two months spanning a possible EE transition (the "source" employer).

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 manifest in Figure 9 and demonstrate that these jumps are not due to unusual seasonality. Specifically, using only the sample of valid SAMEMP answers, we regress the individual EE dummy on dummies for calendar month, rotation group, and on a RIP-based dummy interacted with dummies for each two-month sequence of respondent status (SS, PP, PS, SP, PP'). In the "treatment" specification, the RIP dummy takes value one for cohorts that entered the Survey in 2007 or later, and the regression is run on data on cohorts that entered the survey in either 2006 or 2007. Basically, this regression estimates, for each sequence of respondent status, and controlling for seasonality and rotation group, the difference between the average EE probability of the first cohorts that were exposed to the RIP (2007) and that of the last cohorts that were not (2006). Strictly speaking, this is not a treatment-control specification, because the the control is pre-2007 and the treatment is post-2007, so not simultaneous, and the effect of the RIP may be confounded with other time effects. We are not aware of any other major change in the Survey around that time. Business cycle conditions between 2006 and 2007 are similar: average unemployment rates were 4.6% in those two years and real GDP grew at the relatively healthy pace of 2.9% and 1.9%, respectively, before the Great Recession started to take a heavy toll on the economy in the subsequent years. In the "placebo" specification, the RIP dummy takes value of one for cohorts that entered the Survey in 2006 or later, and the regression is run on data on cohorts that entered the survey in either 2005 or 2006; so, compared with the treatment specification, everything is moved one year back in time, just before the RIP was introduced.

Figure 10 reports the estimated coefficients of the RIP dummy, one per respondent group, and 95% confidence bands. The results indicate that the introduction of the RIP had a sizable, and often statistical significant, negative effect on EE mobility for respondent groups SP, PS, PP, not for SS and PP' (with the caveat that the latter is a small sample). Nothing of the sort was observed a year earlier.

5 Imputation of employer-to-employer transitions

Having provided evidence that the RIP altered measurement of the EE probability, differently by respondent status, we propose an imputation procedure, which exploits data in 1995-2006, before the introduction of the RIP. 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 inhibits many SAMEMP questions. We will also exploit a cyclical indicator of the labor market prospects for each individual, namely, the Unemployment-to-Employment (UE) transition probability for that respondent status group. Both pre-RIP data and random search-and-matching theory predict that UE and EE probabilities should strongly comove over time, due to the common component which is job market tightness, determining the chance of *contact* with open vacancies. This comovement, however, is not perfect, because EE movers are, by virtue of being already employed, pickier and less likely to accept offers.

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 pre- and post-2007 data, given the same observables (worker characteristics, rotation group, respondent status, stage of the business cycle), measures the sample selection of those who do answer the SAMEMP question after 2007, when the RIP is enforced. So, for those who do not answer, 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 out-selection from the sample due to their lower response rates 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 RIP introduced in measuring EE flows, and to obtain a precise imputation formula, we lay out a statistical model. Let E_{it} denote an indicator function that individual i is employed in month t, with observable characteristics Y_{it} (a vector). Recall that $\mathrm{DI}_{it} \in \{0,1\}$ indicates a valid answer to the SAMEMP (Dependent Interviewing, retrospective) question, and let $\mathrm{EE}_{it} \in \{0,1\}$ indicates an employer-to-employer move (that the valid answer is No). A statistical model is

$$
\begin{aligned}
\Pr\left(\mathrm{DI}_{it} = 1 \mid \mathrm{E}_{i,t-1} = \mathrm{E}_{it} = 1\right) &= f_{\mathrm{DI}}(Y_{it}, \theta_{it}) \\
\Pr\left(\mathrm{EE}_{it} = 1 \mid \mathrm{E}_{i,t-1} = 1\right) &= f_{\mathrm{EE}}(Y_{it}, \theta_{it})
\end{aligned}
$$

where θ is an unobservable, whose distribution may depend on observables. We are interested in the average mobility $\mathbb{E}\left[\mathrm{EE}_{it} \mid \mathrm{E}_{i,t-1} = 1\right] = \Pr\left(\mathrm{EE}_{it} = 1 \mid \mathrm{E}_{i,t-1} = 1\right)$ of formerly employed workers for each month t. Some formerly employed workers do not experience an employer-to-employer transition, $\mathrm{EE}_{it} = 0$, because they separate from their job into nonemployment, $\mathrm{E}_{it} = 0$. 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, $\mathbb{E}\left[\mathrm{EE}_{it} \mid \mathrm{E}_{i,t-1} = \mathrm{E}_{it} = 1\right] = \Pr\left(\mathrm{EE}_{it} = 1 \mid \mathrm{E}_{i,t-1} = \mathrm{E}_{it} = 1\right)$ for each month t, but we only observe the realization of their EE_{it} when there is a valid answer $\mathrm{DI}_{it} = 1$, namely $\mathbb{E}\left[\mathrm{EE}_{it} \mid \mathrm{E}_{i,t-1} = \mathrm{E}_{it} = 1, \mathrm{DI}_{it} = 1\right] = \Pr\left(\mathrm{EE}_{it} = 1 \mid \mathrm{E}_{i,t-1} = \mathrm{E}_{it} = 1, \mathrm{DI}_{it} = 1\right)$. The last two expectations do not coincide due to selection by observables and unobservables into giving a valid answer $\mathrm{DI}_{it} = 1$. The unobservable $\theta_{i,t}$ associated to individual i is assumed to be time-varying. Its persistence captures fixed unobserved traits of individual i, such as preference for job stability, that 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 which brings i out of the house for a job interview and triggers a nonresponse.

In principle, we could specify the functions $f_{\mathrm{DI}}, f_{\mathrm{EE}}$ 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 $Y = \mathrm{R} \cup X$: 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 be likely correlated with unobserved heterogeneity. In our empirical implementation, the group R_{it} of record (i, t) will be defined 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 θ_{it}. But, even before the RIP, the R =PP' group exhibits a higher rate of non-response to the SAMEMP question (Figure 7) as well as a higher observed EE probability conditional on valid responses (Figure 9). Therefore, conditioning on group R is useful also before the RIP, as the shares of these respondent groups in the eligible population change over time. Note that, in our specific application, a given individual changes group over time depending on the sequence of respondent status over the last two months. The other observables X_{it} will include surbey rotation group (2-4 and 6-8), demographics of individual i in month t, and a labor market cyclical indicator, the average job finding probability from unemployment of the individual's Respondent group R_{it}.[10]

To ease notation, from now we omit the conditioning on employment in consecutive periods, $\mathrm{E}_{i,t-1} = \mathrm{E}_{it} = 1$, 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:

$$\Pr\left(\mathrm{EE}_{it} = 1 \mid \mathrm{R}_{it}, X_{it}, \theta_{it}\right) = \mathbb{E}\left[\mathrm{EE}_{it} \mid \mathrm{R}_{it}, X_{it}, \theta_{it}\right] = \alpha^{\mathrm{R}_{it}} + X_{it}\beta^{\mathrm{R}_{it}} + \theta_{it} \tag{1}$$

with $\theta \mid R, X \sim G(\cdot \mid \mathrm{R}, X)$ 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

[10]In 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. Defining group by both 5 respondent statuses and 6 rotation groups would require to split the sample each month in 30 groups, which runs into sample size constraints.

R and observables X:

$$\mathbb{E}\left[\mathrm{EE}_{it}\right]=\mathbb{E}_{\mathrm{R},X}\left[\mathbb{E}\left[\mathrm{EE}_{it}\mid \mathrm{R}_{it}=\mathrm{R},X_{it}=X\right]\right] \tag{2}$$

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 which have a valid answer to the SAMEMP question:

$$\begin{aligned}\mathbb{E}\left[\mathrm{EE}_{it}\mid \mathrm{R}_{it},X_{it},\mathrm{DI}_{it}=1\right]&=\mathbb{E}\left[\mathbb{E}\left[\mathrm{EE}_{it}\mid \mathrm{R}_{it},X_{it},\mathrm{DI}_{it}=1,\theta_{it}\right]\mid \mathrm{R}_{it},X_{it},\mathrm{DI}_{it}=1\right]\\&=\mathbb{E}\left[\alpha^{\mathrm{R}_{it}}+X_{it}\beta^{\mathrm{R}_{it}}+\theta_{it}\mid \mathrm{R}_{it},X_{it},\mathrm{DI}_{it}=1\right]\\&=\alpha^{\mathrm{R}_{it}}+X_{it}\beta^{\mathrm{R}_{it}}+\mathbb{E}\left[\theta_{it}\mid \mathrm{R}_{it},X_{it},\mathrm{DI}_{it}=1\right]\end{aligned} \tag{3}$$

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

$$\mathbb{E}\left[\mathrm{EE}_{it}\mid \mathrm{R}_{it},X_{it},\mathrm{DI}_{it}=0\right]=\alpha^{\mathrm{R}_{it}}+X_{it}\beta^{\mathrm{R}_{it}}+\mathbb{E}\left[\theta_{it}\mid \mathrm{R}_{it},X_{it},\mathrm{DI}_{it}=0\right]. \tag{4}$$

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

$$\mathbb{E}\left[\theta_{it}\mid \mathrm{R}_{it},X_{it},\mathrm{DI}_{it}=1\right]\neq\mathbb{E}\left[\theta_{it}\mid \mathrm{R}_{it},X_{it},\mathrm{DI}_{it}=0\right].$$

If this were an equality, we could impute missing records based only on observables, R_{it}, X_{it}, i.e. projecting observed EE_{it} 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 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 θ_{it} 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. $\mathbb{E}\left[\theta_{it}\mid \mathrm{R}_{it},X_{it}\right]=0$.

Given the assumed linear-in-X structure in observables (1), this amounts to assuming that $\mathbb{E}\left[\theta_{it}\mid \mathrm{R}_{it},X_{it}\right]$ is also linear in X, and as such is absorbed in the group fixed effect $\alpha^{\mathrm{R}_{it}}$ and in the term $X\beta^{\mathrm{R}_{it}}$.

Assumption 2: No selection before the RIP. Among respondents who are *not* subject to the RIP (which in particular includes all respondents before the gradual introduction of the RIP that began in January 2007), unobserved heterogeneity θ_{it} is orthogonal to the validity of the answer to the SAMEMP question, conditional on respondent group R_{it} and observables X_{it}:

$$\begin{aligned}\mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1, \mathrm{RIP}_{it}=0\right] = \mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=0, \mathrm{RIP}_{it}=0\right] \\ = \mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}\right] = 0.\end{aligned}$$

This is a MAR (Missing at Random) assumption about answers to the SAMEMP question within each group R_{it} and given other observables X_{it}. 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. Among respondents who are subject to the RIP, mean unobserved heterogeneity amongst valid responses is a time-invariant function $b^{\mathrm{R}}(X)$ of respondent group R and observable characteristics X. For all (i, t):

$$\mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1, \mathrm{RIP}_{it}=1\right] = b^{\mathrm{R}_{it}}(X_{it}).$$

This means that, within each respondent group R, a valid answer to the SAMEMP question when 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 business cycle conditions gathered in the vector X, and has no trend nor other time effects.

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 $\alpha^{\mathrm{R}}, \beta^{\mathrm{R}}$ and $\mathbb{E}\left[\theta \mid \mathrm{R}, X, \mathrm{DI}=0\right]$ for each R, X.

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

$$\begin{aligned}0 = \mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}\right] = \Pr\left(\mathrm{DI}_{it}=0 \mid \mathrm{R}_{it}, X_{it}\right) \cdot \mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=0\right] \\ + \Pr\left(\mathrm{DI}_{it}=1 \mid \mathrm{R}_{it}, X_{it}\right) \cdot \mathbb{E}\left[\theta \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1\right].\end{aligned}$$

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

$$\mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=0\right] = -\frac{\Pr\left(\mathrm{DI}_{it}=1 \mid \mathrm{R}_{it}, X_{it}\right)}{1-\Pr\left(\mathrm{DI}_{it}=1 \mid \mathrm{R}_{it}, X_{it}\right)} \cdot \mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1\right]. \tag{5}$$

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, $\widehat{\mathrm{EE}}_{it}$. Our final time series is the monthly average of these imputed transitions and of observed EE_{it} transitions.

The average "bias" among observed answers, given respondent group and observables, is $\mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1\right]$, which can be decomposed as follows:

$$\begin{aligned}\mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1\right] = {} & \Pr\left(\mathrm{RIP}_{it}=1 \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1\right) \cdot \mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1, \mathrm{RIP}_{it}=1\right] \\ & + \Pr\left(\mathrm{RIP}_{it}=0 \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1\right) \cdot \mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1, \mathrm{RIP}_{it}=0\right].\end{aligned}$$

Now, Assumption 2 implies that $\mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1, \mathrm{RIP}_{it}=0\right]=0$, and Assumption 3 that $\mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1, \mathrm{RIP}_{it}=1\right] = b^{\mathrm{R}_{it}}\left(X_{it}\right)$. Further introducing the notation

$$\phi^{\mathrm{R}}(X) = \Pr(\mathrm{RIP}=1 \mid \mathrm{R}, X, \mathrm{DI}=1)$$

for the conditional probability of a respondent being subject to the RIP, we finally obtain the following expression for the bias:

$$\mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it}=1\right] = \phi^{\mathrm{R}_{it}}\left(X_{it}\right) \cdot b^{\mathrm{R}_{it}}\left(X_{it}\right)$$

The goal of our imputation is to estimate $\phi^{\mathrm{R}}(X)$ and $b^{\mathrm{R}}(X)$ for every R, X.

The key insight behind our imputation strategy is as follows. We know from Table 1 the share of an individual i's cohort that was subject to the RIP in month t, which we denote by $\mathrm{RIPSHARE}_{it} \in \{0, .5, 1\}$ and thus for which cohorts and times $\phi^{\mathrm{R}}(X)=0$ (before January 2007) and $\phi^{\mathrm{R}}(X)=1$ (phased in through March 2009). However, we do not know it for the remaining cohorts (shaded light grey in Table 1), because $\phi^{\mathrm{R}}(X)$ is the probability of being treated by the RIP in the sample of valid SAMEMP answers ($\mathrm{DI}_{it}=1$), while Table 1 only gives us the probability ($\mathrm{RIPSHARE}_{it}=.5$) of being treated by the RIP in the entire sample ($\mathrm{DI}_{it}=0$ or $\mathrm{DI}_{it}=1$) of those cohorts. We can estimate $b^{\mathrm{R}}(X)$ 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 $b(X)$). Then, given Assumption 3, we assume that the estimated $\widehat{b}(X)$ also applies to the cohorts in the RIP transition period (in January 2007-March 2009), and we can use it to estimate $\phi(X)$.

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 ($\mathrm{E}_{i,t-1} = \mathrm{E}_{it} = 1$), every month t run a separate cross-sectional Probit regression of the answer to the SAMEMP question (DI_{it}) on observables X_{it}. Then calculate the predicted value from this regression for each record, and call it $\widehat{P}_{it}$, an estimate of $\Pr(\mathrm{DI}_{it} = 1 \mid \mathrm{R}_{it}, X_{it})$ for that individual.

2. Using all available valid answers to the SAMEMP question ($\mathrm{DI}_{it} = 1$) except for those with $\mathrm{RIPSHARE}_{it} = .5$, run an OLS regression of EE_{it} on a constant, X_{it}, and the interaction of $\mathrm{RIPSHARE}_{it}$ (which in this sample is simply a RIP dummy) with a flexible function $b(X_{it} \mid \gamma)$ parameterized by a vector γ. The resulting estimated coefficients for group R are, resp., $\widehat{\alpha}^{\mathrm{R}}, \widehat{\beta}^{\mathrm{R}}, \widehat{\gamma}^{\mathrm{R}}$. For all records, including those with $\mathrm{RIPSHARE}_{it} = .5$ excluded from the regression, predict $\widehat{B}_{it} = b(X_{it} \mid \widehat{\gamma}^{\mathrm{R}_{it}})$, which estimates the bias of valid answers subject to the RIP (the $b^R(X)$ function introduced in Assumption 3). Let $\widehat{\Phi}_{it} = \mathrm{RIPSHARE}_{it}$, which estimates the share of valid answers who were subject to the RIP. The product, $\widehat{\Phi}_{it} \cdot \widehat{B}_{it}$, estimates the average bias due to the RIP in the observed EE mobility of that cohort.

3. Using all available valid answers to the SAMEMP question ($\mathrm{DI}_{it} = 1$), independently of $\mathrm{RIPSHARE}_{it}$, compute the residual $\widehat{\theta}_{it} = \mathrm{EE}_{it} - \widehat{\alpha}^{\mathrm{R}_{it}} - X_{it}\widehat{\beta}^{\mathrm{R}_{it}}$ based only on the group fixed effect and observables.

4. Using the available valid answers to the SAMEMP question ($\mathrm{DI}_{it} = 1$) *only* in the RIP transition cohorts ($\mathrm{RIPSHARE}_{it}$=.5) regress $\widehat{\theta}_{it}$ on the interaction of $\widehat{B}_{it}$ with a flexible function of observables $\phi(X_{it} \mid \delta)$ parameterized by δ, including a constant term. Then, if the estimate $\hat{\delta}$ coefficients do not exhibit a statistically significant overall pattern, move to the next step; else, only for the sample used in this regression ($\mathrm{DI}_{it} = 1$, $\mathrm{RIPSHARE}_{it}$=.5), update the probability of RIP treatment among valid answers in those cohorts from $\widehat{\Phi}_{it} = \mathrm{RIPSHARE}_{it} = .5$ in Step 2 to $\widehat{\Phi}_{it} = \phi(X_{it} \mid \widehat{\delta}^{\mathrm{R}_{it}})$. Equivalently, update the average bias from $.5\widehat{B}_{it}$ in Step 2 to the fitted value of this regression $\phi(X_{it} \mid \widehat{\delta}^{\mathrm{R}_{it}})\widehat{B}_{it}$.

5. For each eligible record with missing answer $\mathrm{DI}_{it} = 0$, impute

$$\widehat{\mathrm{EE}}_{it} = \widehat{\alpha}^{\mathrm{R}_{it}} + X_{it}\widehat{\beta}^{\mathrm{R}_{it}} - \frac{\widehat{P}_{it}}{1 - \widehat{P}_{it}} \widehat{\Phi}_{it} \cdot \widehat{B}_{it}.$$

6. Every month t, take the sum of EE_{it} when observed and of $\widehat{\mathrm{EE}}_{it}$ when imputed ($\mathrm{DI}_{it} = 0$), 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 ($\mathrm{E}_{i,t-1} = 1$).[11]

By Equation (2), the last ratio is an unbiased estimate of the population average probability of transition from employer to employer. Note that the group of non-eligible records of workers who were formerly employed, but no longer are ($\mathrm{E}_{i,t-1} = 1$, $\mathrm{E}_{it} = 0$), contribute to the denominator ($\mathrm{E}_{i,t-1} = 1$), but are excluded from the numerator, because they would not contribute to it anyway, by $\mathrm{EE}_{it} = 0$. Note also that the imputation is done also for pre-RIP missing records, 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.

We can now more fully illustrate the intuition behind our strategy. The first EE regression in Step 2 exploits Assumption 2 (that pre-RIP records are unbiased because Missing at Random, no selection on unobservables) to compute the $X-$dependent bias post-RIP, $\widehat{B}_{it} = b(X_{it} \mid \widehat{\gamma}^{\mathrm{R}_{it}})$. In Step 4, Assumption 3 ensures that the function $b(X \mid \gamma)$ is time-invariant, so $\widehat{B}_{it} = b(X_{it} \mid \widehat{\gamma}^{\mathrm{R}_{it}})$ applies also in the transition period, to cohorts with $\mathrm{RIPSHARE}_{it} = .5$. Then, among observed answers to the SAMEMP question in those cohorts in transition, we estimate the desired $\mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it} = 1\right]$ not by taking the average of $\widehat{\theta}_{it} = \mathrm{EE}_{it} - \widehat{\alpha}^{\mathrm{R}_{it}} - X_{it}\widehat{\beta}^{\mathrm{R}_{it}}$ for each X_{it} among valid answers, as X contains continuous variables, but as the prediction from the regression of $\widehat{\theta}_{it}$ on $\widehat{B}_{it}$, because $\mathbb{E}\left[\theta_{it} \mid \mathrm{R}_{it}, X_{it}, \mathrm{DI}_{it} = 1\right]$ equals $\phi(X_{it} \mid \delta) b(X_{it} \mid \gamma^{\mathrm{R}_{it}})$. Finally, $\widehat{\Phi}_{it} = \widehat{\phi}(X_{it} \mid \widehat{\delta}^{\mathrm{R}_{it}})$ estimates the incidence of the RIP among respondents in the transition period.

Finally, the smaller the share of missing answers in the survey population, the larger the adjustment in Eq. (5) needed to guarantee that unobserved heterogeneity has zero mean in

[11]For 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 knew 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.

the population, by Assumption 1.[12]

5.3 Imputation regressions: specification and results

In Steps 2 and 4 we specify the functions $b(X_{it} \mid \gamma)$ and $\phi(X_{it} \mid \delta)$ to be linear in the following observables X_{it}: the monthly UE transition probability of the same respondent group, and dummies for: calendar month, rotation group (1-2, 2-3, 3-4, 5-6, 6-7, 7-8), 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, 70-), 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 the nonlinear estimation sometimes fails to converge.

Tables A.1-A.2 report results from the predictive regression in Step 2. We find that the interactions of RIP_{it} with observables X_{it} are often sizable and statistically significant, indicating that the intercept, the cross-sectional relationship between observed EE mobility and demographics−especially age and, to a lesser extent, education−and the time series relationship between EE and UE job finding probability of the same respondent group R changed significantly after the introduction of the RIP. This is further evidence of selection on unobservables. Specifically, the effect of the RIP on average EE mobility controlling for all other covariates, as captured by the RIP coefficient itself (interaction of RIP_{it} with the intercept), is uniformly negative, and significant for some respondent groups. The estimated age profile of EE, which is monotonically decreasing as expected, is weaker post-RIP, suggesting that more mobile individuals are more likely to be missing the younger their age. Similarly, "Never Married" people tend to move more frequently from employer to employer, but less so after the introduction of the RIP. The estimated coefficient on the UE indicator, which is robustly positive as one would expect from its comovement with EE due to the common factor "vacancy postings", is weaker post-RIP, suggesting that the sample gets selected by the RIP in favor of less mobile and cyclically sensitive types.

In Step 4, here is how we gauge the "statistically significant overall pattern" from the results of the regressions of $\widehat{\theta}_{it}$ on the interaction of $\widehat{B}_{it} = b(X_{it} \mid \widehat{\gamma}^{\text{R}_{it}})$ with a flexible function of observables $\phi(X_{it} \mid \delta)$. The results lead to reject at the 1% level the hypothesis that all

[12]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, thus violating Assumption 3. To dispose of this assumption, we can modify Step 4 as follows: regress the estimated residual $\widehat{\theta}_{it}$ on a flexible function of X_{it}, and now define $\widehat{\Phi}_{it}\widehat{B}_{it}$ to be the fitted value from this regression. Since $\widehat{\theta}_{it}$ is built only using $\widehat{\alpha}^{\text{R}}, \widehat{\beta}^{\text{R}}$, which in turn are identified in Step 2 only from pre-2007, pre-RIP data, this strategy also uses less data for imputation.

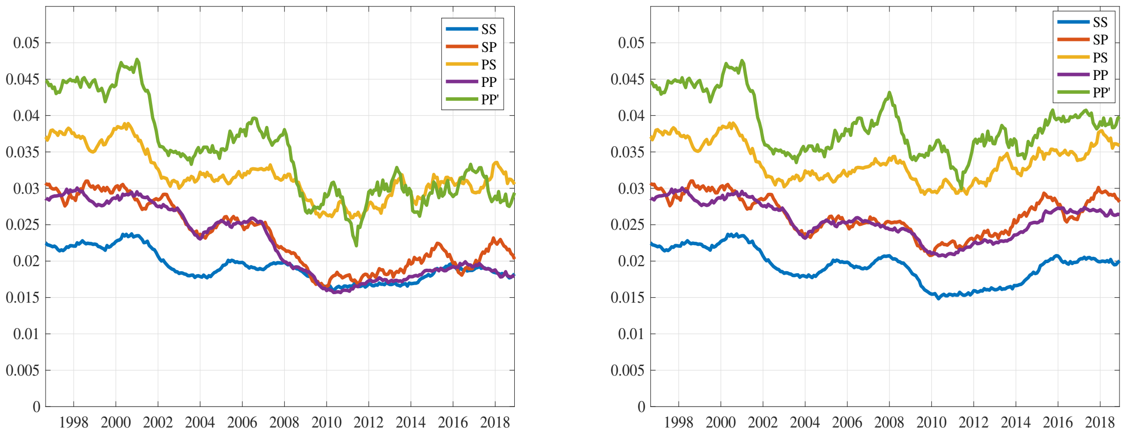

(a) Missing At Random

(b) Imputed

Figure 11: EE Rate by CPS Respondent Group

elements of $\widehat{\delta}$ are jointly statistically insignificant. This is not surprising, because the vector $\widehat{\delta}$ has about 300 components, due to the rich set of observables X at our disposal. Of these 300, however, in each regression, only two or three (depending on the respondent group R) are statistically significant. Moreover, those statistically significant variables are different across regressions, and the signs tend to be different a well. Some of the statistically insignificant elements of $\widehat{\delta}$ are still large numbers, so for these RIPSHARE$_{it}$ = .5 cohorts the fitted values $\phi(X_{it} \mid \widehat{\delta}^{\mathrm{R}_{it}}) \cdot \widehat{B}_{it}$ from these regressions differ somewhat from the covariate $\widehat{B}_{it} = b(X_{it} \mid \widehat{\gamma}^{\mathrm{R}_{it}})$, the baseline "estimated bias" based on Step 2, multiplied by $\widehat{\Phi}_{it}$ =RIPSHARE$_{it}$ = .5. Updating the probability of the RIP treatment among valid answers as described in Step 4, namely replacing $\widehat{\Phi}_{it}$ =RIPSHARE$_{it}$ with $\widehat{\Phi}_{it} = \phi(X_{it} \mid \widehat{\delta}^{\mathrm{R}_{it}})$, makes a small but visible difference to the average EE series especially in 2007, when the upward correction is even stronger, but also in 2008. For all these reasons, we prefer a conservative approach and do not use the results of Step 4 regression, but keep the baseline $\widehat{\Phi}_{it}$ =RIPSHARE$_{it}$ for all records.

5.4 Final results: the imputed EE series

In Figure 11, we report, for each survey respondent group R=SS,PS,SP,PP,PP' on which we perform the imputation separately, the time series for the average monthly employer-to-employer mobility rate in the US since 1995, estimated using the Missing at Random assumption (MAR), and our imputation method. All time series are MA-smoothed to remove

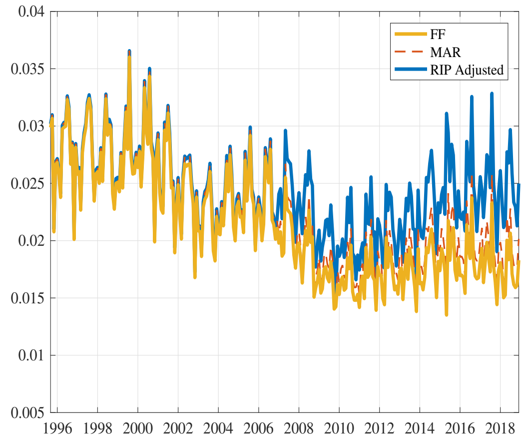

(a) Raw Series

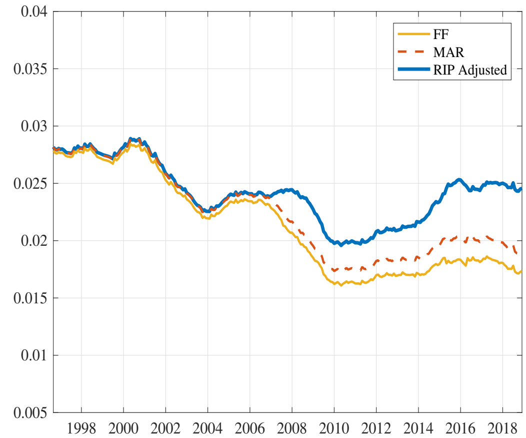

(b) 12-Month Trailing Moving Average

Figure 12: EE Rate: Fallick and Fleischman (2004) vs. Missing At Random vs. Imputed

high frequency noise.

The imputed series, which start diverging from the raw ones after the introduction of the RIP in 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 Coding questions, including SAMEMP, exhibit observable characteristics that strongly correlate in other records with EE mobility.

In Figure 12, 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 employer-to-employer transition in the US since 1995, estimated using the Fallick and Fleischman (1994) method (FF), the Missing at Random assumption (MAR), and our imputation method. In the second panel, all time series are MA-smoothed to remove high frequency noise.

By an unfortunate coincidence, the introduction of the RIP preceded by about a year the onset of the Great Recession. Since EE probabilities are procyclical, the sharp drop observed around 2008 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 and due to the RIP. While the imputed EE probability did fall, importantly, it declined later, and by much less than the raw series, as well as than the UE transition probability, which declined by about half starting in late 2008, following the financial crisis. The imputed and FF/MAR series share a weak recovery in 2010-2014, and then a clear rebound, which ends in 2016. Thereafter, our EE series remains flat at the pre-Great Recession level of about 2.5%, while the raw series remain 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 the EE measure of turnover appears to have stabilized.

5.5 Comparison with the LEHD

To further corroborate the validity of our imputation, in Figure 13, we compare the three time series from Figure 12 with a fourth one, derived from a different, administrative, matched employer-employee dataset, the Longitudinal Household and Employer Dynamics (LEHD). This source does not suffer from missing answers, but still requires an imputation, because EE flows are computed from quarterly reports on total earnings accruing to each worker from each employer over the entire previous quarter, thus are subject to time aggregation bias. Specifically, we know when a worker earned income from one employer A in a quarter, from two employers, A and B, in the following quarter, and only from B in the quarter after that. To label this an EE transition, however, we have to rule out 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.[13]

To facilitate comparison with the seasonally adjusted quarterly series of the LEHD-based employer-to-employer transition rate, available from 2000:Q2 to 2017:Q4, we seasonally adjust (Census X12) our monthly CPS series, take quarterly averages, and rescale them so that the average level of our imputed/MAR series match up with the average of the LEHD series for the first three years of the sample. The correlation coefficient of each of our three series with the LEHD series is 0.7491 (FF), 0.8012 (MAR), 0.9314 (our imputed series). The stronger correlation of our series is close to that shared by all measures pre-RIP (2000:Q2-2006:Q4): 0.8857(FF), 0.8888 (MAR), 0.8930 (our imputed series). Just like our imputed series, the LEHD series starts dropping visibly with the onset of the recession, in early 2008, not in late 2006 like the existing FF and MAR estimates based on the RIP, which is further evidence of the effect of the RIP.

The only visible difference between our imputed series and the LEHD is that, during the

[13] `https://lehd.ces.census.gov/data/j2j_beta.html`.

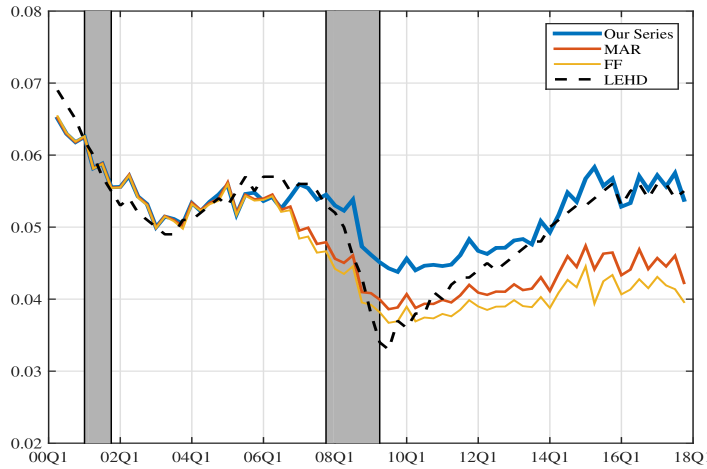

Figure 13: Quarterly EE probability: CPS and LEHD

Note: Shaded areas indicate NBER dated recessions.

Great Recession, the LEHD drops proportionally a lot more than our imputed series, and indeed than any CPS series. This difference raises the concern that our imputation maybe over-correcting the drop due to the RIP. We believe, however, that this large drop in LEHD is due to time aggregation. The described procedure to eliminate 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 last 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.

6 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, coinciding with the introduction of new interviewing policy, the Respondent Identification Policy (RIP). We provide evidence that these answers are not missing at random, and the RIP caused a serious 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 transition rate. We show that the decline in the aggregate employer-to-employer monthly transition probability 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.

Our analysis still faces an important limitation. The share of invalid answers to the SAMEMP question in the CPS was modest but slowly rising even before the introduction of RIP in 2007; after that episode, the share jumps up, but then continues rising 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 unobservable is unfolding, unrelated to RIP and partially immune to our imputation. In future research, we plan to investigate the causes of these trends, and especially their sharp acceleration caused by the introduction of the RIP. Getting to the bottom of this measurement issue is especially important in light of the recent debate on declining dynamism in US labor market.

References

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, 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 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.

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

Appendix

A.1 Imputation Step 2: regression results

Tables A.1 and A.2 summarize the results of the Step-2 regressions in our imputation procedure, as described on page 27. Table A.1 presents the coefficient estimates for the pre-RIP sample ($\text{RIP}_{it} = 0$) together with the effect of the RIP_{it} dummy itself. The results are largely plausible. For example, older age groups tend to be associated with lower EE transitions and UE rates are positively related to EE transitions.

The coefficients on the RIP_{it} dummy tend to be negative (see the first row), although only two of them are statistically significant. As shown in Table A.2, however, many of the RIP_{it} interaction terms are highly statistically significant. As described in Assumption 3 on page 25, those with valid answers ($\text{DI}_{it} = 1$) are allowed to be systematically less (or more) mobile than those with missing answers ($\text{DI}_{it} = 0$) and this differential mobility can depend on demographics and business cycle conditions. Our results indicates that this is indeed the case.

A.2 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_{it} 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 five respondent group (as in our proposed imputation procedure) over three different samples: $\text{RIPSHARE} = \{0, 0.5, 1\}$. 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.

In Figure A.1, 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.

Table A.1: Step 2 Regressions (I)

	$R = SS$	$R = SP$	$R = PS$	$R = PP$	$R = PP'$
RIP	-0.004	-0.007	-0.012^{*}	-0.004^{*}	0.013
	(-1.476)	(-1.101)	(-1.708)	(-1.876)	(1.360)
Rotation Gr. 2−3	-0.002^{***}	-0.001	-0.008^{***}	-0.003^{***}	-0.004^{**}
	(-6.166)	(-1.374)	(-7.544)	(-6.419)	(-2.305)
Rotation Gr. 3−4	-0.003^{***}	-0.002^{*}	-0.012^{***}	-0.003^{***}	-0.009^{***}
	(-8.786)	(-1.948)	(-12.455)	(-6.455)	(-4.782)
Rotation Gr. 5−6	-0.000	0.003^{***}	-0.001	-0.000	-0.000
	(-0.184)	(2.804)	(-0.472)	(-0.377)	(-0.161)
Rotation Gr. 6−7	-0.002^{***}	-0.001	-0.009^{***}	-0.003^{***}	-0.005^{***}
	(-4.920)	(-1.049)	(-8.914)	(-6.022)	(-3.003)
Rotation Gr. 7−8	-0.004^{***}	-0.001	-0.012^{***}	-0.003^{***}	-0.006^{***}
	(-11.014)	(-1.579)	(-11.411)	(-7.318)	(-3.230)
Sex	0.002^{***}	-0.002^{***}	0.000	0.001^{*}	-0.001
	(6.249)	(-2.998)	(0.728)	(1.819)	(-1.051)
Married Spouse Absent	0.004^{***}	0.007^{***}	0.020^{***}	0.010^{***}	0.010^{***}
	(7.386)	(3.218)	(7.740)	(8.935)	(3.154)
Widowed/Divorced	0.003^{***}	0.004^{***}	0.014^{***}	0.008^{***}	0.010^{***}
	(9.232)	(4.274)	(11.023)	(13.390)	(4.673)
Never Married	0.003^{***}	0.008^{***}	0.021^{***}	0.009^{***}	0.004^{**}
	(8.530)	(8.343)	(18.301)	(19.710)	(2.130)
High School	-0.003^{***}	-0.003^{**}	-0.007^{***}	0.001^{***}	0.007^{***}
	(-5.686)	(-2.432)	(-5.637)	(2.777)	(4.727)
Some College	-0.001	-0.002	-0.007^{***}	0.005^{***}	0.010^{***}
	(-1.618)	(-1.335)	(-5.447)	(8.623)	(6.325)
College	-0.001	-0.001	-0.006^{***}	0.003^{***}	0.013^{***}
	(-0.981)	(-1.142)	(-4.070)	(5.544)	(5.923)
Graduate	0.000	0.001	-0.006^{***}	0.003^{***}	$0.005*$
	(0.756)	(0.626)	(-3.872)	(4.949)	(1.745)
Ages 21−30	-0.025^{***}	-0.023^{***}	-0.016^{***}	-0.016^{***}	-0.016^{***}
	(-16.052)	(-9.962)	(-6.598)	(-21.067)	(-8.718)
Ages 31−40	-0.033^{***}	-0.036^{***}	-0.027^{***}	-0.026^{***}	-0.028^{***}
	(-21.177)	(-15.108)	(-10.790)	(-32.471)	(-11.972)
Ages 41−50	-0.036^{***}	-0.038^{***}	-0.030^{***}	-0.029^{***}	-0.039^{***}
	(-23.321)	(-16.053)	(-11.697)	(-35.478)	(-16.207)
Ages 51−60	-0.038^{***}	-0.037^{***}	-0.033^{***}	-0.026^{***}	-0.041^{***}
	(-24.279)	(-15.298)	(-12.873)	(-31.205)	(-16.445)
Ages 61−70	-0.039^{***}	-0.036^{***}	-0.032^{***}	-0.024^{***}	-0.039^{***}
	(-24.674)	(-13.880)	(-11.747)	(-25.112)	(-11.223)
Ages 71−	-0.040^{***}	-0.034^{***}	-0.033^{***}	-0.021^{***}	-0.040^{***}
	(-23.064)	(-9.708)	(-9.275)	(-13.532)	(-6.036)
UE	0.052^{***}	0.038^{***}	0.024^{***}	0.052^{***}	0.042^{***}
	(15.308)	(6.838)	(4.386)	(14.137)	(4.355)
Constant	0.037^{***}	0.046^{***}	0.054^{***}	0.029^{***}	0.025^{***}
	(17.057)	(11.529)	(12.519)	(15.838)	(3.805)

Table A.2: Step 2 Regressions (II): Interaction Terms with RIP_{it}

	$R = SS$	$R = SP$	$R = PS$	$R = PP$	$R = PP'$
Rotation Gr. 2−3	−0.000	0.001	−0.001	0.001	−0.004
	(−0.457)	(0.794)	(−0.826)	(1.582)	(−1.372)
Rotation Gr. 3−4	−0.001	0.002	0.001	0.000	0.005*
	(−1.030)	(1.340)	(0.808)	(0.056)	(1.936)
Rotation Gr. 5−6	0.001**	0.001	−0.001	0.001**	0.000
	(2.291)	(0.952)	(−0.337)	(2.277)	(0.064)
Rotation Gr. 6−7	−0.001	0.000	−0.001	0.001**	−0.001
	(−1.008)	(0.200)	(−0.416)	(2.410)	(−0.282)
Rotation Gr. 7−8	0.000	−0.000	−0.000	0.001	0.002
	(0.440)	(−0.253)	(−0.007)	(0.955)	(0.896)
Sex	−0.001*	0.000	0.001	−0.001***	0.002
	(−1.736)	(0.314)	(1.347)	(−3.137)	(0.937)
Married Spouse Absent	−0.001*	−0.004	−0.001	−0.005***	−0.008*
	(−1.708)	(−1.120)	(−0.329)	(−3.826)	(−1.906)
Widowed/Divorced	−0.000	−0.003*	−0.002	−0.005***	−0.011***
	(−1.037)	(−1.812)	(−0.827)	(−6.917)	(−3.380)
Never Married	0.000	−0.004***	−0.007***	−0.004***	−0.002
	(0.134)	(−3.100)	(−4.184)	(−6.744)	(−0.868)
High School	−0.000	0.002	0.008***	−0.000	−0.002
	(−0.327)	(0.872)	(3.715)	(−0.388)	(−1.031)
Some College	−0.000	0.000	0.008***	−0.001*	−0.004
	(−0.456)	(0.256)	(3.691)	(−1.794)	(−1.583)
College	−0.000	0.001	0.010***	−0.000	−0.006*
	(−0.190)	(0.565)	(4.417)	(−0.223)	(−1.800)
Graduate	−0.000	−0.002	0.011***	−0.001	0.001
	(−0.332)	(−0.799)	(4.426)	(−0.910)	(0.236)
Ages 21−30	0.013***	0.006*	0.005	0.006***	0.008***
	(5.707)	(1.650)	(1.070)	(5.511)	(2.753)
Ages 31−40	0.015***	0.013***	0.0086*	0.009***	0.012***
	(6.886)	(3.415)	(1.762)	(8.540)	(3.329)
Ages 41−50	0.017***	0.012***	0.0106**	0.011***	0.018***
	(7.485)	(3.144)	(2.309)	(9.804)	(4.883)
Ages 51−60	0.017***	0.009**	0.012**	0.007***	0.019***
	(7.677)	(2.248)	(2.541)	(6.143)	(4.979)
Ages 61−70	0.018***	0.006	0.012**	0.004***	0.018***
	(7.872)	(1.485)	(2.434)	(3.514)	(3.651)
Ages 71−	0.020***	0.005	0.015**	0.003	0.014
	(8.124)	(0.933)	(2.550)	(1.363)	(1.573)
UE	−0.028***	−0.022***	−0.008	−0.030***	−0.037***
	(−6.634)	(−2.825)	(−0.931)	(−6.333)	(−2.637)
R^2	0.004	0.007	0.008	0.007	0.010
Sample Size	4421008	861278	857044	3870909	318530

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 UE rate for each respondent group enters the regression for the corresponding group. The sample period is September 1995 - December 2018 but excludes the cohorts with RIPSHARE = 0.5 (See Table 1). The superscripts *, **, and *** indicate significance at 10%, 5%, and 1%, respectively. t-statistics are in parentheses.

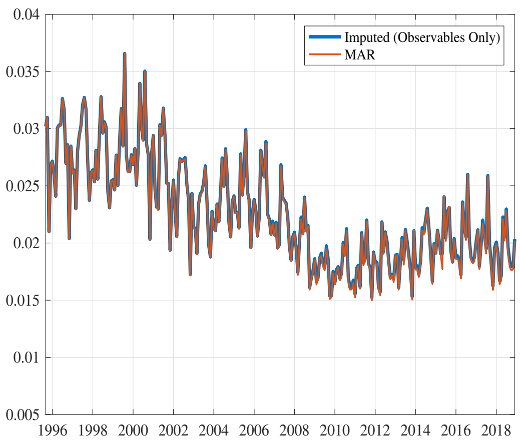

(a) Raw Series

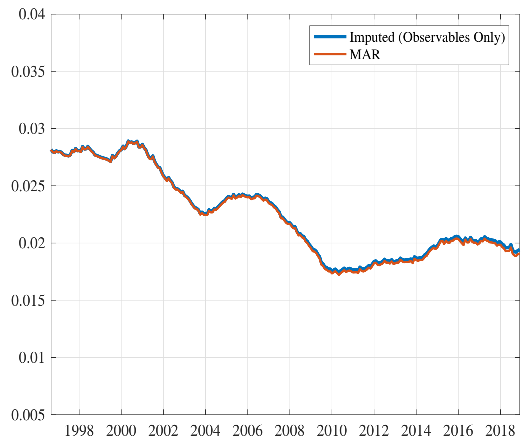

(b) 12-Month Trailing Moving Average

Figure A.1: EE Rate: Missing At Random vs. Imputed by Observables Only