

Measuring Employer-to-Employer Reallocation

By SHIGERU FUJITA, GIUSEPPE MOSCARINI, AND FABIEN POSTEL-VINAY*

We revisit measurement of Employer-to-Employer (EE) transitions in the monthly Current Population Survey. The incidence of missing answers to the question on change of employer sharply increases starting with the introduction of a new software instrument to conduct interviews in January 2007 and of the Respondent Identification Policy in 2008-2009. We document non-random non-response selection by observable and unobservable worker characteristics that correlate with EE mobility. We propose a selection model and a procedure to impute missing answers. Our imputed EE aggregate series no longer trends down after 2000 and restores a close congruence with the business cycle after 2007.

JEL: J63, E24

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, researchers have been paying increasing attention to the flow of workers from Employer to Employer (EE), with no intervening jobless spell. A prominent literature, as best exemplified by Burdett and Mortensen (1998) and its empirical applications, as well as by Postel-Vinay and Robin (2002), shows that on-the-job search by, and competition between firms for, employed workers are a natural source of worker bargaining power, and an important determinant of cross-sectional wage dispersion caused by frictional turnover.

Just as critical is the role that EE reallocation plays in shaping two dynamic aspects of US labor markets. First, from the individual point of view of a typical US worker, direct moves from one employer to another are a major source of earnings growth over the life cycle (Topel and Ward, 1992), but also of idiosyncratic earnings risk. Climbing the job ladder takes time; therefore, falling off it can have drastic implications for lifetime earnings (Davis and Von Wachter (2011), Jarosch (2023), Huckfeldt (2022)) and explain the striking skewness and kurto-

* Fujita: Federal Reserve Bank of Philadelphia, Research Department, Ten Independence Mall, Philadelphia, PA 19106, shigeru.fujita@phil.frb.org. Moscarini: Yale University, Department of Economics, Cowles Foundation for Research in Economics, and National Bureau of Economic Research, 87 Trumbull Street, New Haven, CT 06511, giuseppe.moscarini@yale.edu. Postel-Vinay: University College London, Department of Economics and Institute for Fiscal Studies, Drayton House, 30 Gordon Street, London WC1H 0AX, United Kingdom, f.postel-vinay@ucl.ac.uk. We thank the Co-Editor, Aysegul Sahin, and two anonymous referees for extensive and useful comments that greatly enriched the contents of this paper; Anne Polivka for her assistance on various technical aspects of the Current Population Survey; Henry Hyatt (discussant), Antoine Bertheau (discussant), Jim Spletzer, and participants at numerous conferences and seminars, for comments and discussions. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of Philadelphia or any other person associated with the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

sis in individual earnings growth at annual frequency documented by Guvenen, Ozkan and Song (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 either Unemployment or Non-participation into Employment. A large share of these UE and NE flows comprises, respectively, recalls by the last employer (Fujita and Moscarini, 2017) and first-time entry into the labor force, which do not directly reallocate workers between firms. Therefore, EE transitions play a quantitatively dominant role in this type of reallocation, a major driver of aggregate productivity growth (e.g. Foster, Haltiwanger and Syverson, 2008; Lentz and Mortensen, 2008). The EE transition probability is also procyclical, but much less volatile than the UE probability or the unemployment rate. These facts bear significant implications for the cyclical reallocation of labor input between firms, industries, and occupations (Haltiwanger et al., 2018), for the estimation of the matching function (Moscarini and Postel-Vinay, 2018), and for measurement of mismatch and labor market slack relevant to monetary policy (Moscarini and Postel-Vinay, 2023).

For all these reasons — and possibly more — measuring EE transitions accurately is important. This is the goal of the present paper.¹ We focus on the monthly Current Population Survey (CPS, Census Bureau (2022a)), the premier source of real-time information on labor markets, including the civilian unemployment rate, available to policymakers in the United States. The monthly frequency, almost unique among labor force surveys even in developed countries, reduces both recall bias and the time aggregation that blurs the distinction between direct EE transitions and short unemployment spells in quarterly survey data. Moscarini and Postel-Vinay (2016) obtain an EE series in 1996-2013 from the Survey of Income and Program Participation (SIPP); the results of the SIPP are published with significant delay, and its quality deteriorated since 2014, when the interview frequency declined from thrice to once a year. The Job Openings and Labor Turnover Survey (JOLTS, Bureau of Labor Statistics (2022)) publishes its latest measurements a month after the monthly CPS, including a Quits series, which conflates quits to other jobs with quits to non-employment, two radically different outcomes. Quits have been rising exceptionally fast during the recovery from the pandemic recession, and are at an all-time high, giving rise to a “Great Resignation” narrative. In contrast, our imputed EE series peaked in mid-2021, and has been stable-to-falling ever since. Finally, administrative data, most notably the quarterly Longitudinal Employer Household Dynamics (LEHD) dataset (Census Bureau (2022b)), provide a measure of EE transitions free of

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

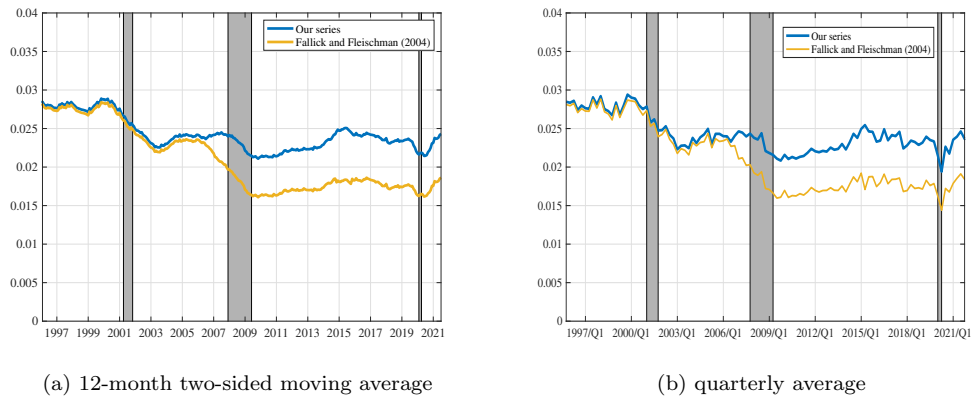


Figure 1. : Employer-to-Employer (EE) transition probability

Note: Shaded areas indicate NBER dated recessions (NBER (2023)).

survey measurement error, but suffer from severe time aggregation (see Section V for details), lack of unemployment measures, and long processing times. We conclude that the CPS remains the benchmark.

In the late 1980s and early 1990s, the Census identified unacceptable measurement error in a number of individual transitions (between employers, industries, occupations, employment status) constructed by comparing answers given by CPS respondents in consecutive months to the same question, such as employer name, or job tasks. Small changes in verbal descriptions resulted in an enormous number of spurious transitions, detected by administrative data for some subsamples. To address this issue, in the 1994 Survey redesign, the Census introduced Dependent Interviewing. This consists of explicit retrospective questions: in the case of employer changes, the interviewer reads out the name of an individual’s employer recorded in the previous month, and simply asks if it still the same, a yes/no answer (recorded in the variable IODP1). While the name of the employer may remain incorrectly coded, the transition is more accurately detected. In this paper, following the IPUMS terminology (IPUMS CPS (2023)), we will refer to this question as “EMPSAME.”² Fallick and Fleischman (2004) pioneered its use to estimate the average EE monthly transition probability, a time series that became the standard reference in the profession (Fallick and Fleischman (2019)). The lighter line in Figure 1a shows the time series of our replication of their results, after taking a 12-month two-sided Moving Average to eliminate high-frequency noise; Figure 1b shows quarterly averages of the seasonally-adjusted series. We can see a dramatic decline that starts in 2006, and never reverts, thus generating

²Moscarini and Thomsson (2007) show that the introduction of Dependent Interviewing in 1994 suddenly reduced measured industry and occupational transitions by 90%, and exploit post-1994 cleaner data to impute the pre-1994 observations.




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

the impression of a sharp cyclical drop preceding the Great Recession by a full year, as well as of a downward trend, and a similarly dramatic but transient drop in April-May 2020, during the COVID-19 crisis.

In this paper, we revisit measurement of the EE transition probability. Our starting point is Figure 2. We detect a sudden and sharp increase in the incidence of missing answers to the EMPSAME question, starting in January 2007, followed by a further acceleration through 2009, which never reversed and continued growing gradually through 2015. We identify one important change in survey methodology phased in starting in January 2008 by the US Census Bureau, the Respondent Identification Policy (RIP), which directly impacts the validity of the answer to the EMPSAME question. In a nutshell, the RIP gives, for privacy reasons, the respondent the option not to share their answers, including employer names that make the EMPSAME question possible, with any other household members who might happen to answer the survey in subsequent months. A significant number of respondents exercise that option, automatically generating a missing answer to the EMPSAME question a month later. This sudden censoring generates a very strong sample selection on unobservable worker characteristics, which correlate with EE mobility. That is, given a group of observationally identical individuals who are eligible to activate the RIP, once some of them do, and suddenly stop answering the EMPSAME question, the average EE mobility of the remaining valid answers in that group simultaneously, and just as suddenly, drops. We also detect another source of measurement error of different nature than the RIP, affecting all CPS cohorts in 2007, coinciding in time with the switch by the Census to a new software instrument to conduct interviews.³ For all these

³In addition, as reported in this IPUMS webpage (IPUMS CPS (2023)), the employer name was not

reasons, observed EE transitions after 2007 poorly estimate the true incidence of EE reallocation. Because the EMPSAME question was introduced only to correctly detect industry and occupational transitions, the RIP directly impacts also these two important series.⁴

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 EMPSAME question, thus EE transitions, both before and especially after January 2007. The issue we face is a classic one of self-selection into treatment. The basic idea behind our imputation is simple. Conditional on very detailed observable worker characteristics, including the change in the identity of the Survey respondent within the household that triggers the application of the RIP, and on the state of the aggregate labor market, we estimate the jump in the share of missing answers to the EMPSAME question and the jump in measured EE probability around the time of the RIP introduction, 2008-2009. Then, we use the observed jump in missing EMPSAME responses to apportion back the missing EE transitions, conditional on observables. That is, we attribute the jumps in both series to the RIP, but allow its effects to vary flexibly by worker observables, so that the jumps reflect selection by unobservables, conditional on observables. To project our imputation forward in time to 2010-2022, we make the key identification assumption that this RIP-induced selection bias is time-invariant, conditional on the worker characteristics and aggregate (trend and business cycle) indicators used in the imputation. We apply a similar, but separate, procedure to 2007 observations, impacted only by the new interviewing software.

By implementing our procedure, we estimate an aggregate EE time series which differs substantially, over the last 15 years, from Fallick and Fleischman's (2019), plotted as a dark line in Figure 1.⁵ Specifically, our series resets the cyclical peak from 2006 to 2008, more in line with evidence from administrative data reviewed later, and reduces the subsequent cyclical drop by about half, with a full recovery by 2015, followed by a mild decline thereafter. Thus, our imputed series restores a closer congruence between EE transitions and the business cycle, greatly reduces their cyclical volatility, and eliminates the appearance of a "quit-less recovery" after the Great Recession and of declining EE dynamism in the US labor market since the early 2000s. Interestingly, at the monthly frequency, both EE and (as already emphasized by Crump et al. (2019) with quarterly data) UE inflows into

recorded in the first rotation group in May 2015, for unknown reasons, so all EMPSAME answers are missing a month later. This is why we drop the observation for that month in Figure 2.

⁴In the online Appendix, we examine industry and occupational mobility, as well as four more CPS variables that utilize Dependent Interviewing, thus are potentially affected by the RIP: self-employment, retirement, disability, and unemployment duration. We show that, unlike EE and industry/occupational mobility, none of these four are materially impacted by the RIP, either because Dependent Interviewing applied differently, or because the Census could rely on other cross-sectional questions, immune to the RIP, to measure those other states.

⁵We make available at Fujita, Moscarini and Postel-Vinay (n.d.), and regularly update, the EE time series that we estimate based on both Fallick and Fleischman's (2019) and our methodology, as well as the one based on a Missing at Random assumption. See also Figure 14(a), which plots all three time series.

employment show no appreciable trend in the XXI century, while the EU outflow probability (see also Fujita (2018)) and measures of job reallocation and business entry kept trending down. This robust set of facts should introduce important nuances in the debate on declining fluidity of the US labor market.

We also present the first empirical evidence of a large, albeit transient, negative impact of the COVID-19 crisis on the pace of EE reallocation; our imputed EE probability series drops even more dramatically than the Fallick and Fleischman (2004)'s series, as predicted by the selection model, because response rates to the EMPSAME questions rose sharply, reflecting the observed higher availability of previous survey respondents under home lockdown.

Finally, we relate our revised measure of EE mobility in the US, based on the monthly CPS, to the time series of the same transition probability estimated from the other three mentioned datasets: LEHD, SIPP, and JOLTS. Despite their limitations, that we highlighted, we find the comparison useful. We reconcile findings across datasets, which generate cyclically synchronized series, albeit with different volatilities in LEHD and JOLTS, and support the delayed cyclical decline in 2008 and no trend in EE mobility in the CPS in this century.

The paper is organized as follows. In Section I we illustrate the features of the monthly CPS designed to detect individual EE transitions, with a detailed description of the pertinent EMPSAME question. In Section II 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 changes in the interview instrument and protocol, most notably the RIP in 2008. In Section III we provide evidence that the RIP significantly changed measured EE transitions. In Section IV we propose and implement our imputation procedure of missing answers, hence of EE transitions, based on a model of selection by unobservable worker characteristics that affect the propensity both to answer the survey and to change job. In Section V we compare our imputation results with those from other datasets, we address the impact of survey attrition, and we examine the impact of the COVID-19 pandemic. Brief conclusions take stock of the results and highlight open issues in the measurement of labor market transitions in the CPS, that we leave for future research.

I EE transitions: data and definitions

A *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 us to detect employer changes has been available only since the 1994 survey redesign, as described below.⁶ Despite

⁶Most of the overview information presented in this section is directly based on the official description of the CPS (Census Bureau (2022a)).

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

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

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

As is well-known, estimating transitions in the monthly CPS requires matching records, namely, uniquely identifying records in consecutive survey months that refer to the same individual. 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 matching records in post-1993 data. Appendix A.A1 illustrates our matching algorithm. In general, matching probabilities are fairly high, although over the past several years attrition grew by about two percentage points.

B The 1994 survey redesign: Dependent Interviewing

An overhaul of the interviewing technique took place in 1994.⁷ Among many other changes, one was critical to our exercise. Until 1993, *every* month, respondents were asked anew: (i) for whom they worked, (ii) what kind of business that was, (iii) what kind of work they were doing, (iv) what their most important activities were, and (v) what sector they were working in. This information was later used by CPS staff to assign employer, occupation and industry codes to each individual. This “Independent Coding” procedure had at least two serious shortcomings. First, asking these questions was very cumbersome for the interviewer, and respondents typically complained about answering the same questions repeatedly. Second, and more important for our purposes, asking these questions independently every month introduced a significant amount of spurious shifts in occupation and industry. Indeed, in a small validation study of occupational coding based on company records and employees’ descriptions of their own tasks, Mathiowetz (1992) finds that CPS staff coded occupations incorrectly about half the time when not told that two consecutive records concern the same individual. More remarkably, when told that the two records *did* come from the same individual, these expert coders still found a 12% disagreement rates between the company record and the employee’s description of their tasks.

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

- 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 additional questions ask about occupation in the new employer, which is then coded independently of the previous one. If the answer is Yes, then the respondent is only asked to confirm the description of activities given a month before, either by themselves or by another respondent in the

⁷This description is based on Polivka and Rothgeb (1993). See also Moscarini and Thomsson (2007).

household. In that case, Dependent Coding applies and automatically assigns the same occupational code as in the previous month.

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

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

II Missing answers to the EMPSAME question

A Facts

Among CPS records matched between months $t - 1$ and t , those who are employed in both months are eligible for the EMPSAME question in month t . Throughout the paper, whenever we mention “eligibility,” we refer to this criterion and, unless otherwise explicitly stated, analyze this eligible sample. In this sample, we count those who answer No to this question. The ratio between this count and the total number of employed in the initial month within the matched sample is our measure of the EE probability.⁸

The highest hurdle in this apparently straightforward computation is caused by missing answers to the EMPSAME question among eligible records. Those

⁸Note that the denominator includes some individuals who are no longer employed in the current month.




Figure 3. : EE^m : Missing answers to the EMPSAME question in eligible (employed in both months) records

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 they are at the same company or not. The question is whether the true, unobserved answer was positive or negative. The issue is real even for small percentages of missing answers, because the raw monthly EE probability, computed by just discarding records with missing answers, is small (around 2%), and we do not know the conditional EE probability among those missing answers. For example, suppose that only 1% of all answers are missing but that they are all EE movers in truth. Then, the true EE monthly probability would increase by one half, from 2% to roughly 3%, adding on average close to 20 million EE transitions per year in the US.

In Figure 3, the higher, darker line illustrates how the share of eligible records with missing answers to the EMPSAME question (EE_t^m) evolved since the introduction of Dependent Interviewing. Four facts stand out. First, this share has always been positive and non-negligible. Second, it has been rising over time. Both facts were already noticed by Moscarini and Thomsson (2007), their 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 concern. There is further sharp acceleration through early 2009. Fourth, the share drops visibly in April and May 2020 and then rebounds, following the COVID-19 shock.

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

8) at the beginning of 2009, explaining the sharp accelerations in the aggregate measure.

Fallick and Fleischman (2004) pioneered the use of Dependent Interviewing to calculate this EE probability, and their time series has become the main reference in the profession. We reconstruct their time series (Fallick and Fleischman (2019)), using their described methodology and assuming, as they do, that missing answers to the EMPSAME question are stayers.⁹ Our reverse-engineered time series and the one that Fallick and Fleischman make available on their website (Fallick and Fleischman (2019)) coincide almost perfectly, as we show in Figure 4(a), where the two lines lie on top of each other.¹⁰ Their treatment of missing answers is potentially problematic even before 2007, more so since then. Instead of treating missing answers as stayers, one can assume that the EE probability of these missing answers is the same as that among valid responses. In Figure 4(b), we can see that this Missing-At-Random (MAR) assumption brings the level of the EE probability up noticeably. The gap has been widening since around 2007, in line with increasing incidence of non-response shown in Figure 3.

The natural question is: what happened in 2007-2009? Adams et al. (2019) contains, on p.37, a chronology of changes in monthly CPS interviewing protocols. Only two changes are mentioned in that time frame. First, in January 2007 the system running the data collection instrument changed from a DOS-based system to Blaise, a Windows-based system. Quoting from p.103 of the same document: “The instrument consists of complicated skip patterns and automated question text fills.” Skipping patterns are essential to Dependent Interviewing questions, such as EMPSAME. Second, starting in January 2008 the Census phased in the Respondent Identification Policy (RIP). We now explain the nature of the RIP and provide evidence that these two changes in interviewing caused a sudden, uneven, and temporary increase in measurement error in 2007 (Blaise), and a more severe, gradual one starting in 2008, which became permanent in 2009 (RIP). While the software change in 2007 applies to all interviews, the RIP only applies to some types of respondents, that we are able to identify.

B 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

⁹This assumption is not described in Fallick and Fleischman (2004), but was confirmed in a private communication with Charles Fleischman, whom we thank.

¹⁰Fallick and Fleischman (2004) also exclude Rotation Groups 1 and 5 from their calculations, to avoid the so-called “first rotation group bias,” and focus on transitions between months in sample 2-3, 3-4, 6-7, and 7-8. We follow them to replicate their series in Figure 4. In the rest of our analysis, however, we include all Rotation Groups, including 1 and 5, thus transitions between months in sample 1-2 and 5-6, because we find that they make little difference to the aggregate time series, but they increase the sample size for our imputation procedure of missing answers to the EMPSAME question, described later.

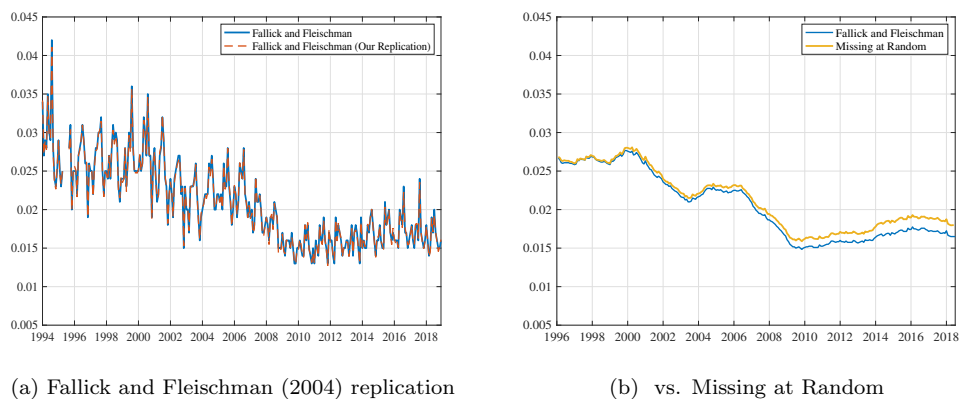


Figure 4. : EE probability: Fallick and Fleischman (2004) Series

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

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.

Census Bureau (2015) describes the RIP in Chapter 2.D. “If the original respondent, which we refer to as the ‘RIP respondent,’ wishes their information to be confidential, and they are not available for a subsequent interview, you cannot conduct dependent interviewing. However, if the RIP respondent permits you to verify their information with anyone in the household, then you can conduct dependent interviewing. [...] The instrument will only allow one person to be the RIP respondent. Once the RIP question is asked and the RIP respondent is selected, the RIP question will not be re-asked in subsequent months. You may change the answer to the RIPFLG question during the initial interview only. The only time the RIPFLG will change in subsequent interviews is when there is a replacement household.” Therefore, once the RIP is implemented, a negative answer by the first, RIP respondent invalidates dependent interviewing for the entire 4+4 month sequence of that household, unless the household moves out of the address and is replaced by another one moving in. For this reason, although “Any household member 15 years of age or older is technically eligible to act as a

respondent”, the Manual then continues: “If at all possible, try to interview the most knowledgeable member of the household. In most situations, this individual will be the reference person or the spouse of the reference person.” In turn, the reference person is defined as “The first person mentioned by the respondent, who either owns or rents the ‘sample unit’ (e.g., house, apartment).”

Polivka et al. (2009) also report that the RIP question is not asked in single-person households, while 14.4% of the RIP questions that were asked for all of 2008 received a negative answer, from respondents who are observationally different from the population. One concern for our purposes is that employed and job-mobile respondents are more likely to answer No to the RIP question, suggesting that they have some confidentiality concerns about their work situation, primarily about their earnings. Polivka et al. (2009) also report that, in 2008, following one of the 14.4% negative answers to the RIP questions, the respondent changed in only one in nine (11%) households in the following month’s interview. Multiplying the two shares, the No response to the RIP question should result in a share of invalid dependent interviewing of just about 1.5%. We showed much larger numbers than this, *especially* after 2008, because a No answer to the initial RIP question has ramifications that propagate to all other household members and beyond the month of the answer and the one following it, and suppresses information.

From now on, our strategy proceeds in three steps.

First, in the remainder of this section, we estimate the timing and mode of introduction of the RIP in monthly interviews. The variable RIPFLG flags when an interview is subject to the RIP, and contains the answer to the RIP question for that household, but is not available in the public use data, nor in any confidential version of the data that we are aware of. To determine when and how the RIP was rolled out, we thus proceed indirectly. Based on Polivka et al. (2009) and Adams et al. (2019), we estimate that the RIP was introduced in 2008, but gradually. To validate this prior, we exploit the fact that the RIP invalidates some answers to the EMPSAME question. Then, we measure the occurrence and size of month-over-month changes in the share of missing answers to the EMPSAME question, EE_t^m , starting in 2006. We do this for each cohort and rotation group. The RIP applies if both the household has more than one member and the household member who answers from the second month on differs from the original RIP respondent. Therefore, we dig deeper into the pattern of EE^m missing answers to the EMPSAME question among eligible records, breaking it down by household size (one/more) and by respondent status (Self/Proxy). Consistently with our assumption, respondent groups that we expect to be more affected by the RIP show the largest jumps in EE_t^m , in calendar months that we identify.

Our second step, in Section III, exploits the exogenous variation across groups in the timing of the RIP introduction, to identify whether the RIP, or something else, caused a change in measured EE transitions. It is highly unlikely that other changes, especially in the labor market, affected those rotation groups exactly

in those months and in that same order. Thus, we use a “treatment-control” approach to document that, every time the RIP was rolled out for a group of respondents, i.e., the share of valid answers to the EMPSAME question suddenly declined, so did the measured EE probability among the remaining valid answers, only for that specific rotation group. So changes in the incidence of EE_t^m cause simultaneous drastic changes in measured EE, which is the object of interest. We also provide other, auxiliary evidence to support our hypothesis that the switch to the Blaise software in January 2007 and the phase-in of the RIP starting in January 2008 are the only causes of the sudden increase in missing answers to the EMPSAME question.

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

C Identification of survey respondents

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

For this purpose, we use variable PUSLPRX that indicates whether the person answered the survey that month for the household, to identify the respondent (PULINENO) for each household (HRHHID and HRHHID2). We then construct a flag taking values SS, SP, PS, PP, and PP', and we assign it to each of the matched records between the two adjacent months. That is, the answer to the EMPSAME question in month t is flagged, say, PS if that answer was given by a Proxy in month $t - 1$ and by the individual them-Self in month t . Single-person households are easily identified and necessarily belong to the SS group. To the best of our knowledge, we are the first to construct such a variable in the monthly CPS and to study its implications for the measurement of labor market transitions.

Figure 5 plots the shares of the five groups in the population of eligible (employed both last and this month) matched records in each calendar month. The shares of SP and PS are virtually identical. We can see that the share of each




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

group is roughly constant until around 2007, and then SS and PP start rising, presumably reflecting the Census Bureau's effort to secure the same respondent in consecutive interviews after the roll-out of the Blaise software in 2007 and of the RIP in 2008. All shares exhibit sharp temporary blips at the onset of the pandemic in Spring 2020.

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

Figure 6 plots the shares of EE_t^m by respondent status, namely, the proportion of eligible records *within each respondent group* that has no valid answer to the EMPSAME question. These shares rise over time in each group. Consistently with the logic of the RIP, since 2008, these shares are lower (more valid answers) when the respondent's identity does not change (SS, PP) and higher when it changes and the person in question responds neither time (PP'). Indeed, these shares suddenly rise in January 2007 for all respondent groups, and again very fast in early 2008 only for the SP, PS and PP' groups, while SS and PP show no unusual behavior after 2007. This evidence is consistent with the change in the




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

CAI software system introduced in January 2007 affecting *all* respondent groups, and with the RIP introduction in January 2008 affecting *only* the three groups with respondent turnover. The CAI effect does not disappear in 2008 and beyond, because the missing answers in the SS, PP groups, to which the RIP does not apply, never revert to pre-2007 levels.

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

The natural experiment of the recent COVID-19 crisis offers additional evidence in support of our hypothesis that the RIP caused the non-response rate to the EMPSAME question to rise since 2008. In Figure 5, we can see that respondent turnover drops drastically in April and May 2020: the shares of SS and PP, which are immune to the RIP, rise by about eight percentage points. Presumably, forced by the lockdown to stay at home, people were suddenly more available to respond again to the CPS. Accordingly, in Figure 7, the share of missing answers to the




Figure 7. : Missing answers to the EMPSAME question among Self Responses

EMPSAME question falls suddenly and dramatically, briefly returning to levels not seen in about a decade. Some of this drop, however, occurs even within Respondent groups, conditional on the type of respondent turnover, as seen in Figure 6. This is a change in behavior, not in selection, which tempers the impact of the RIP. We can only speculate on the reasons, as also probably related to cohabitation during the lockdowns.

D Timing of RIP roll-out

Let $RIP_{i,t} \in \{0, 1\}$ indicate whether the RIP applies to the survey respondent who answers questions regarding individual i in month t ,¹¹ and $DI_{i,t} \in \{0, 1\}$ whether a valid answer to the EMPSAME (Dependent Interviewing, retrospective) question regarding individual i in month t is available. Note that i refers to the identity of the person who is the subject of the questions, not to the identity of the respondent. If $RIP_{i,t} = 1$, so the survey asks the RIP question, and the answer is No, then the EMPSAME question cannot be asked and $DI_{j,s} = 0$ for all members $j \neq i$ of the same household and all calendar months $s \geq t$ when the household is interviewed. But it is also possible that the EMPSAME question can be asked and yet the respondent refuses to answer, or does not know the answer, in which case we have $DI_{i,t} = 0$.

Let $\Pr(DI_{i,t} = 0)$ denote the probability of an invalid answer to the EMPSAME question among eligible records in month t , which can be estimated by the observed share EE_t^m of invalid answers. Note that $DI_{i,t}$ is an individual-

¹¹That is, $RIP_{i,t}=1$ whenever individual i at time t is part of a household whose first rotation respondent (not necessarily i) was asked the RIP question and gave an answer stored in the Census variable RIPFLG.

level variable, while EE_t^m is an aggregate time series, a population share, whose time series is plotted in Figure 3. Let $\Pr(\text{RIP}_{i,t} = 1)$ be the probability of a record in month t being subject to the RIP. While we do not observe $\text{RIP}_{i,t}$, we are extremely confident that $\Pr(\text{RIP}_{i,t} = 1) = 0$ before January 2008 and $\Pr(\text{RIP}_{i,t} = 1) = 1$ starting sometime in 2009, given the evidence in Figures 1b and 3 and the description in Polivka et al. (2009). Then we estimate before 2008

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

and after 2009

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

To estimate the object of interest, $\Pr(\text{RIP}_{i,t} = 1)$, in the 2008-2009 period, we use the identity:

$$\begin{aligned} \Pr(\text{DI}_{i,t} = 0) &= \Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t} = 1) \Pr(\text{RIP}_{i,t} = 1) \\ &\quad + \Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t} = 0) \Pr(\text{RIP}_{i,t} = 0) \end{aligned}$$

and make the following identification assumption: $\Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t})$ is constant over time for either $\text{RIP}_{i,t} = 0$ or 1 in a period of time surrounding the RIP roll-out, 2006-2010, so we can estimate $\Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t} = 0)$ for t in the roll-out period 2008-2009 with the average of $\Pr(\text{DI}_{i,\tau} = 0) = EE_\tau^m$ in months $\tau \in 2006$, and $\Pr(\text{DI}_{i,t} = 0 \mid \text{RIP}_{i,t} = 1)$ with the average of $\Pr(\text{DI}_{i,\tau} = 0) = EE_\tau^m$ in months $\tau \in 2010$. Then, using our estimate $\Pr(\text{DI}_{i,t} = 0) = EE_t^m$, the last equation can be solved to obtain an estimate of the incidence of the RIP in every month t of its roll-out period between January 2008 and December 2009:

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

In words, we assume that the entire increase in the incidence of missing answers to the EMPSAME question in this interim 2008-2009 period is due to the introduction of the RIP, and is proportional to the share of records introduced to the RIP. We perform this estimation for each rotation group separately.

To further refine our estimate of the interim period, we zoom onto the period surrounding 2004-2013, and add two more pieces of information. First, we examine the time series of $\Pr(\text{DI}_{i,t} = 0)$ for single-member household, who are not subject to the RIP, and thus are never asked that question. For this special group, Figure 8 shows a jump in January 2007, which reverses in February 2008. Therefore, calendar year 2007 is not affected by the RIP, but by the new Blaise software. Also, after 2007, there is no trend. This is in contrast to the average population, hence to multi-member households, who are vulnerable to




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

the RIP, so their rising trend in non-response rate must be related to the RIP. Second, we break down the time series of EE_t^m not only by respondent status (SS, PS etc.), as done in Figure 6, but also by rotation group. To save on space, in Figure 9 we only present results for the SS and PS group, as PP and SP are (respectively) similar. The SS group (top row) shows modest upward jumps in the incidence of missing answers in January 2007, possibly reversed in early 2008, as for single-person households in Figure 8, due to “growing pains” in using the new Blaise software. Indeed, the rise is most pronounced for Rotation Groups 1 and 5, to which the RIP does not apply. Rotation Groups 5-7 also show a jump in early 2009, staggered in order of rotation (RG5 jumps first, then RG6 a month later, etc.). Conversely, the PS group (bottom row), more likely than both SS and single-person households to be affected by the RIP, shows small jumps in January 2007 and huge jumps, again upwards, in January of 2008 for RG1-3 and 2009 for RG5-7, again staggered in order of rotation.

We conclude that the RIP was introduced in a staggered manner, by rotation group, starting in January 2008, while during the entire 2007 calendar year the transition to Blaise software affected *all* records.

Table 1 provides an overview of our estimated timing of the RIP roll-out period. The date in each cell represents the survey start month (cohort) and the first column gives the calendar time. All cohorts and rotation groups are subject to the transition to Blaise software that causes a temporary increase in the incidence of missing Dependent Interviewing answers in the calendar year 2007, indicated by dates in italics in the middle block. The RIP is introduced by CPS cohort, starting with the one that entered the survey in January 2008. From that point on, all new cohorts are exposed to the RIP (boldfaced dates in the lower block).



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

The RIP roll-out is completed in April 2009, when the last cohort not exposed to it (December 2007) exits the survey.¹²

Figure 10 plots the same EE_t^m share series as in Figure 3, but with respect to the *cohort* dates, i.e. the dates when each cohort entered the survey that are the entries in the table, rather than with respect to the *calendar* dates. We can clearly see a large jump in the January 2008 cohort, as well as a jump in late 2005 followed by gradual increases toward January 2007. This pattern is consistent with Table 1. The oldest cohort that is newly interviewed with Blaise software in January 2007, in their last month in sample, is the October-2005 cohort (right upper corner of Table 1, so only one-eighth of that cohort was subject to that error (only in

¹²Our imputation procedure of EE mobility after 2007 will provide additional evidence of the pattern illustrated in Table 1. The “bias” introduced by the RIP in measured EE, that we estimate for each respondent group and that we aim to correct, settles into a perfectly regular seasonal pattern after 2008, while it is more erratic in 2007. See Figure A3.

Table 1—: Timing of changes in interviewing methodology in the monthly CPS, 2007-2009.

Calendar date	Rotation Group							
	1	2	3	4	5	6	7	8
2006-01	2006-01	2005-12	2005-11	2005-10	2005-01	2005-12	2004-11	2004-10
2006-02	2006-02	2006-01	2005-12	2005-11	2005-02	2005-01	2004-12	2004-11
2006-03	2006-03	2006-02	2006-01	2005-12	2005-03	2005-02	2005-01	2004-12
2006-04	2006-04	2006-03	2006-02	2006-01	2005-04	2005-03	2005-02	2005-01
2006-05	2006-05	2006-04	2006-03	2006-02	2005-05	2005-04	2005-03	2005-02
2006-06	2006-06	2006-05	2006-04	2006-03	2005-06	2005-05	2005-04	2005-03
2006-07	2006-07	2006-06	2006-05	2006-04	2005-07	2005-06	2005-05	2005-04
2006-08	2006-08	2006-07	2006-06	2006-05	2005-08	2005-07	2005-06	2005-05
2006-09	2006-09	2006-08	2006-07	2006-06	2005-09	2005-08	2005-07	2005-06
2006-10	2006-10	2006-09	2006-08	2006-07	2005-10	2005-09	2005-08	2005-07
2006-11	2006-11	2006-10	2006-09	2006-08	2005-11	2005-10	2005-09	2005-08
2006-12	2006-12	2006-11	2006-10	2006-09	2005-12	2005-11	2005-10	2005-09
2007-01	<i>2007-01</i>	<i>2006-12</i>	<i>2006-11</i>	<i>2006-10</i>	<i>2006-01</i>	<i>2005-12</i>	<i>2005-11</i>	<i>2005-10</i>
2007-02	<i>2007-01</i>	<i>2007-12</i>	<i>2006-12</i>	<i>2006-11</i>	<i>2006-02</i>	<i>2006-01</i>	<i>2005-12</i>	<i>2005-11</i>
2007-03	<i>2007-03</i>	<i>2007-02</i>	<i>2007-01</i>	<i>2006-12</i>	<i>2006-03</i>	<i>2006-02</i>	<i>2006-01</i>	<i>2005-12</i>
2007-04	<i>2007-04</i>	<i>2007-03</i>	<i>2007-02</i>	<i>2007-01</i>	<i>2006-04</i>	<i>2006-03</i>	<i>2006-02</i>	<i>2006-01</i>
2007-05	<i>2007-05</i>	<i>2007-04</i>	<i>2007-03</i>	<i>2007-02</i>	<i>2006-05</i>	<i>2006-04</i>	<i>2006-03</i>	<i>2006-02</i>
2007-06	<i>2007-06</i>	<i>2007-05</i>	<i>2007-04</i>	<i>2007-03</i>	<i>2006-06</i>	<i>2006-05</i>	<i>2006-04</i>	<i>2006-03</i>
2007-07	<i>2007-07</i>	<i>2007-06</i>	<i>2007-05</i>	<i>2007-04</i>	<i>2006-07</i>	<i>2006-06</i>	<i>2006-05</i>	<i>2006-04</i>
2007-08	<i>2007-08</i>	<i>2007-07</i>	<i>2007-06</i>	<i>2007-05</i>	<i>2006-08</i>	<i>2006-07</i>	<i>2006-06</i>	<i>2006-05</i>
2007-09	<i>2007-09</i>	<i>2007-08</i>	<i>2007-07</i>	<i>2007-06</i>	<i>2006-09</i>	<i>2006-08</i>	<i>2006-07</i>	<i>2006-06</i>
2007-10	<i>2007-10</i>	<i>2007-09</i>	<i>2007-08</i>	<i>2007-07</i>	<i>2006-10</i>	<i>2006-09</i>	<i>2006-08</i>	<i>2006-07</i>
2007-11	<i>2007-11</i>	<i>2007-10</i>	<i>2007-09</i>	<i>2007-08</i>	<i>2006-11</i>	<i>2006-10</i>	<i>2006-09</i>	<i>2006-08</i>
2007-12	<i>2007-12</i>	<i>2007-11</i>	<i>2007-10</i>	<i>2007-09</i>	<i>2006-12</i>	<i>2006-11</i>	<i>2006-10</i>	<i>2006-09</i>
2008-01	2008-01	<i>2007-12</i>	<i>2007-11</i>	<i>2007-10</i>	<i>2007-01</i>	<i>2006-12</i>	<i>2006-11</i>	<i>2006-10</i>
2008-02	2008-02	2008-01	<i>2007-12</i>	<i>2007-11</i>	<i>2007-02</i>	<i>2007-01</i>	<i>2006-12</i>	<i>2006-11</i>
2008-03	2008-03	2008-02	2008-01	<i>2007-12</i>	<i>2007-03</i>	<i>2007-02</i>	<i>2007-01</i>	<i>2006-12</i>
2008-04	2008-04	2008-03	2008-02	2008-01	<i>2007-04</i>	<i>2007-03</i>	<i>2007-02</i>	<i>2007-01</i>
2008-05	2008-05	2008-04	2008-03	2008-02	<i>2007-05</i>	<i>2007-04</i>	<i>2007-03</i>	<i>2007-02</i>
2008-06	2008-06	2008-05	2008-04	2008-03	<i>2007-06</i>	<i>2007-05</i>	<i>2007-04</i>	<i>2007-03</i>
2008-07	2008-07	2008-06	2008-05	2008-04	<i>2007-07</i>	<i>2007-06</i>	<i>2007-05</i>	<i>2007-04</i>
2008-08	2008-08	2008-07	2008-06	2008-05	<i>2007-08</i>	<i>2007-07</i>	<i>2007-06</i>	<i>2007-05</i>
2008-09	2008-09	2008-08	2008-07	2008-06	<i>2007-09</i>	<i>2007-08</i>	<i>2007-07</i>	<i>2007-06</i>
2008-10	2008-10	2008-09	2008-08	2008-07	<i>2007-10</i>	<i>2007-09</i>	<i>2007-08</i>	<i>2007-07</i>
2008-11	2008-11	2008-10	2008-09	2008-08	<i>2007-11</i>	<i>2007-10</i>	<i>2007-09</i>	<i>2007-08</i>
2008-12	2008-12	2008-11	2008-10	2008-09	<i>2007-12</i>	<i>2007-11</i>	<i>2007-10</i>	<i>2007-09</i>
2009-01	2009-01	2008-12	2008-11	2008-10	2008-01	<i>2007-12</i>	<i>2007-11</i>	<i>2007-10</i>
2009-02	2009-02	2009-01	2008-12	2008-11	2008-02	2008-01	<i>2007-12</i>	<i>2007-11</i>
2009-03	2009-03	2009-02	2009-01	2008-12	2008-03	2008-02	2008-01	<i>2007-12</i>
2009-04	2009-04	2009-03	2009-02	2009-01	2008-04	2008-03	2008-02	2008-01

Note: The date within each cell indicates the survey start month (cohort date). Dates in italic (middle block) indicate that survey respondents in the cohort are subject to measurement error in Dependent Interviewing due to the transition to Blaise software. Dates in bold (bottom block) indicate that respondents in the cohort are subject to the RIP at that date.

their last rotation). The November-2005 cohort had two interviews subject to that error; the December-2005 cohort had three interviews..., the January-2007 cohort had all eight interviews, and this remains the case for all cohorts until December-2007 included. So when we plot EE^m by cohort (as in Figure 10), it rises only gradually from October 2005 through January 2007. After that, it remains roughly constant during 2007, until the January-2008 cohort, when the RIP is introduced to that cohort and subsequent ones for all eight rotations, with a much more dramatic impact on EE^m .

III 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 EMPSAME question and provide three pieces of empirical evidence that the RIP introduced a strong selection.

First, Figure 13(a) later in the paper plots the average EE probability of each





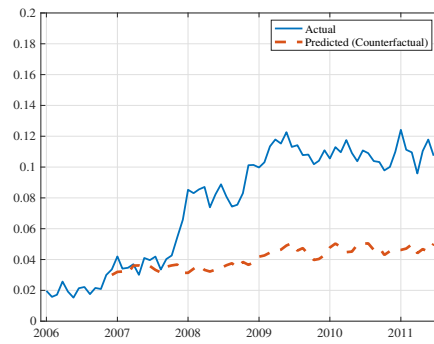
Figure 10. : Missing answers to the EMPSAME question by CPS cohort (month of entry into the CPS)

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 EMPSAME question). EE probabilities differ very significantly across groups, so the changing composition by group of valid answers to the EMPSAME question, documented earlier, affects in itself the aggregate EE probability. More importantly, now PS and SP are no longer equivalent. The former has a much higher EE probability than SP, which is instead similar to PP. Note that these two-month respondent groups only seldom include the initial, RIP respondent, so PP may be affected by the RIP if the first respondent was S (or a different proxy P').


Second, these differences are due to unobservable individual characteristics that correlate with respondent status, rather than to a different composition of each respondent group by observables such as demographics. In Step 1 of our imputation algorithm, described in the next section, we run a Probit regression of the $DI_{i,t}$ dummy on a rich set of observables, separately for each month (starting in 2007) and respondent group: SS, PS, SP, PP, PP'. The estimated coefficients of the regressions based on the SS sample, the group largely immune to the RIP, are often statistically significantly different from the estimated coefficients for the other groups. Using the SS coefficients to impute the probability of missing answers ($DI_{i,t}=0$) to the other groups (PS, SP, PP, PP'), we can produce counterfactual time series of the average incidence of missing answers for those groups, as if they had the same characteristics as SS. As shown in Figure 11, when compared to the actual incidence of missing answers in the data (solid line), the counterfactual incidence (dashed line) is slightly higher in 2007 and, for the SP, PS and PP' groups that are most affected by the RIP, much higher starting in January 2008.




(a) SS




(b) SP



(c) PS



(d) PP



(e) PP'


Figure 11. : Composition effects in missing EMPSAME answers

In each case, the counterfactual incidence shows, after 2007, a mild upward trend that aligns almost perfectly with the actual trend in missing answers pre-2007. This evidence supports our hypothesis that the jump in 2008 is entirely related to the identity of the respondents, in a manner that can only be rationalized by the application of the RIP, and not at all to their observable characteristics.


Finally, we run “treatment/control” and “placebo” experiments, which quantify the jumps that are visually manifest in EE probabilities for the five respondent groups (see Figure 13(a) presented later in the paper). Specifically, using only the sample of valid EMPSAME answers from January 2006 to March 2009, we regress the individual $EE_{i,t}$ dummy on dummies for month of the year and rotation group, and on two treatment dummies, which mark the middle and lower blocks in Table 1. These two treatment dummies are interacted with respondent status dummies (SS, PP, PS, SP, PP’). The first treatment dummy, that we will denote by $BLAISE_{i,t}$, equals one if an observation is in the middle block (dates in italics) of Table 1, which flags the measurement problem related to the Blaise software, and zero otherwise; and the second treatment dummy, $RIP_{i,t}$, equals one if an observation in the lower block (dates in boldface) of Table 1, which indicates exposure to the RIP, and zero otherwise. The observations in 2006 are in the control group and subject to neither of the measurement problems. This regression estimates, after controlling for seasonality and the rotation group, the differences in average EE probabilities of the two treatment groups relative to that of the control group. For the “placebo” experiment, we take the 2005-2006 sample and “treat” 2006 observations with a RIP placebo. That is, the (placebo) RIP dummy takes zero for the 2005 sample and one for the 2006 sample. Basically, we estimate the correlation between absence of answers to the EMPSAME question and negative answers among the valid ones.

Panels (a) and (b) of Figure 12 summarize the impacts of the two dummies. Both of those dummies are associated with significantly lower EE probabilities for all respondent types except SS. For the first, Blaise treatment, the largest impact is observed among PP, which one can also notice in the time series presented below in Figure 13(a). The RIP treatment results in further declines in EE probabilities. Interestingly, EE probabilities among SS are little affected by either treatment. Our placebo regression (Panel (c)) shows no indication of similar declines a year earlier.


Note that, strictly speaking, these are not treatment/control regressions, because the experiments are not simultaneous and therefore the effects of the treatments can be confounded with other time effects. In particular, the treatment periods include the Great Recession, which officially started in December 2007. This is particularly problematic for the second dummy (i.e., the RIP dummy) which marks observations in 2008 and early 2009 (see lower block of boldfaced dates in Table 1). For a genuine treatment/control regression we focus on the period between January 2008 and March 2009. One (randomly-selected) half of this sample is subject only to the change in Blaise software, and the rest are sub-




(a) Treatment 1 (Blaise software, year 2007)



(b) Treatment 2 (RIP, lower left block in Table 1)



(c) Placebo (2006 vs 2005)



(d) Treatment 2 (RIP) - Treatment 1 (Blaise)

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

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

ject to the RIP. The treatment dummy equals one for those exposed to the RIP and zero otherwise. For the first calendar month of the sample (January 2008), only the January 2008 cohort (the first Rotation Group) is subject to the RIP, and remaining cohorts are not. For the last calendar month in the sample (March 2009), all rotation groups except the last one (December 2007 cohort) are subject to the RIP. This sample structure allows us to identify the effect of the RIP (in addition to the month effect and the rotation group effect) relative to that of the 2007 Blaise software change, controlling for the time effect. Panel (d) of Figure 12 presents the estimated coefficients on the RIP dummy. We can see that the RIP tends to be associated with lower (measured) EE probabilities, particularly among SP and PS, although the effect on the PP' group is imprecisely estimated and, in this case, not statistically significant at the 95% confidence level.

IV Imputation of EE transitions

We provided empirical evidence that changes in interviewing protocols in the CPS, conducted with new software since January 2007 and under the RIP since January 2008, altered measurement of EE probabilities, differently by respondent status. To redress measurement, we propose an imputation procedure based on data from 1995-2006. A simple approach is to impute EE assuming no selection by unobservable worker characteristics. This will be approximately correct only if observable worker characteristics strongly correlate with the unobservable ones that determine both true EE mobility and the valid answer to the EMPSAME 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 EMPSAME question, and the two-month respondent status sequence (SS, PP, PS, SP, PP'). Regarding the latter, more job-mobile individuals/households may be more likely to trigger a change in respondent status (SP, PS, PP') and thus the application of the RIP, which prevents the interviewer from asking the EMPSAME question. We will also exploit an aggregate indicator of labor market prospects for each individual to capture common factors, both trend and business cycle, that affects everybody's true EE probability, independently of the RIP.

If sizable unobserved heterogeneity remains after conditioning on observables, the resulting imputation will not correct for the entire bias in the raw series. Therefore, we introduce a model of selection on unobservables. The difference in average EE probabilities between pre- and post-2007 data, given the same observables (worker characteristics, rotation group, respondent status, aggregate indicator), measures the sample selection of those who do answer the EMPSAME question after 2007. So, for those who do not answer the EMPSAME question, the missing records that we want to impute, the bias is the opposite of this difference, scaled by proportions of valid and invalid records. For example, if individuals who are more affected by the RIP tend to have a *higher* true EE probability, then their selection out of the sample will make the bias in the post-RIP observed EE probability negative, more so the larger the relative incidence of missing records. Our imputation model formalizes this insight.

A Imputation: model

In order to clarify the possible sources of bias that the changes in the CPS interviewing protocol, especially the RIP, introduced in measuring EE flows, and to obtain a precise imputation formula, we lay out a statistical model. When we implement this procedure, we treat observations from middle block (italicized dates) of cohorts in Table 1 (Blaise treatment) separately.

Let $E_{i,t}$ denote an indicator function that individual i is employed in month t , with observable characteristics $Y_{i,t}$ (a vector). Recall that $DI_{i,t} \in \{0, 1\}$ indicates

a valid answer to the EMPSAME retrospective question, and $EE_{i,t} \in \{0, 1\}$ an employer-to-employer move (that the answer to EMPSAME is a valid No). A statistical model is

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

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

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

We partition observables Y into two sets $Y = R \cup X$: a "group" R that will be treated nonparametrically, namely, imputation will be performed for each set of individuals in each group separately; and a vector X that will enter parametrically, through regressions using data within each group R . The variables defining the R partition should be likely to be correlated with unobserved heterogeneity. In our empirical implementation, we define a group R by respondent status (SS,SP,PS,PP,PP'), which triggers application of the RIP, which in turn may invalidate eligible records for reasons possibly related to unobserved heterogeneity $\theta_{i,t}$. But, even before 2007, the $R = PP'$ group exhibits a higher rate of non-response to the EMPSAME question (Figure 6) as well as a higher observed EE probability conditional on valid responses (Figure 13(a)). Therefore, condi-

tioning on respondent group $R \in (\text{SS}, \text{SP}, \text{PS}, \text{PP}, \text{PP}')$ is useful also before 2007 as the shares of these respondent groups in the eligible population change over time.¹³ Note that, in our specific application, a given individual changes respondent group over time depending on the sequence of respondent status over the last two months. The other observables $X_{i,t}$ are discussed below.

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

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

(1)

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

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

Our goal is to estimate the average EE transition rate in the population. By the L.I.E., we can write it as the EE probability conditional on respondent groups R and observables X , averaged over these conditioning variables:

(2)

$$\mathbb{E}[\text{EE}_{i,t}] = \mathbb{E}_{R,X}[\mathbb{E}[\text{EE}_{i,t} \mid R_{i,t} = R, X_{i,t} = X]]$$

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 EMPSAME question:

(3)

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

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

(4)

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

¹³In principle, the individual's unobserved propensity to change job is also correlated with rotation group, because people who move to a different address to take a new job are no longer present in later rotation groups, the well-known issue of geographical attrition in the CPS. In Section V, comparing with other datasets, we show evidence that survey attrition is quantitatively a minor concern for EE measurement. Defining group by both 5 respondent statuses and 6 rotation group pairs (1-2, 2-3, 3-4, 5-6, 6-7, 7-8) requires splitting the sample each month in 30 groups, which runs into sample size constraints.

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

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

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

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

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

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

Assumption 2: No selection before 2007. Among respondents to the EMPSAME question who are *not* subject to the interviewing software change and to the RIP, unobserved heterogeneity $\theta_{i,t}$ is orthogonal to the validity of the answer to the EMPSAME question, conditional on respondent group $R_{i,t}$ and observables $X_{i,t}$:

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

This is a MAR (Missing at Random) assumption about answers to the EMPSAME question within each group $R_{i,t}$ and given other observables $X_{i,t}$. Therefore, before 2007, missing responses to the EMPSAME question are immune from selection on unobservables.

Assumption 3: Time-invariant selection after 2007 conditional on observables. For records subject to the RIP, mean unobserved heterogeneity amongst valid responses to the EMPSAME question is a time-invariant

function $b^R(X)$ of respondent group R and observable characteristics X . For all (i, t) :

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

A similar assumption applies to records collected using new Blaise software in 2007.

Assumption 3 implies that, within each respondent group R , a valid answer to the EMPSAME question when the respondent is exposed to the RIP may indicate a systematically higher (or lower) mobility than a valid answer to EMPSAME when *not* exposed to the RIP, but this differential mobility only depends on demographics and aggregate labor market conditions gathered in the vector X , and has no residual trend nor other time effects. Because we treat the Blaise software source of measurement error and the RIP separately, this assumption applies to either, each with its own time-invariant function. Note that, while we assume a time-invariant bias *function*, the actual bias can change over time for observationally identical individuals, because X can contain observable time effects such as trends and business-cycle indicators.

B Imputation: implementation

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

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

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

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

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

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-2007. We can then use those estimates in Equation (4) to impute to each missing record an estimated probability of an employer-to-employer move, $\widehat{EE}_{i,t}$. Our final time series is the monthly average of these imputed transitions and of observed $EE_{i,t}$ transitions.

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

$$\begin{aligned} & \mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1] \\ &= \Pr(\text{RIP}_{i,t} = 1 \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1) \cdot \mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1, \text{RIP}_{i,t} = 1] \\ &+ \Pr(\text{RIP}_{i,t} = 0 \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1) \cdot \mathbb{E}[\theta_{i,t} \mid R_{i,t}, X_{i,t}, DI_{i,t} = 1, \text{RIP}_{i,t} = 0]. \end{aligned}$$

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

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

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

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

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

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

- 3) For each eligible record i, t with missing answer $DI_{i,t} = 0$, impute

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

- 4) Every month t , take the sum of $EE_{i,t}$ when observed ($DI_{i,t} = 1$) and of $\widehat{EE}_{i,t}$ when imputed ($DI_{i,t} = 0$) across all eligible records, so across all respondent groups R and observables X , and divide it by the number of matched individuals in the same CPS cohort who were employed a month before ($E_{i,t-1} = 1$).¹⁴

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

We can illustrate the intuition behind our strategy. The first EE regression in Step 2 exploits Assumption 2 (pre-RIP records are unbiased because Missing at Random and of no selection on unobservables) to compute the X -dependent bias post-RIP, $\widehat{B}_{i,t} = b(X_{i,t} | \widehat{\gamma}^{R_{i,t}})$. This bias summarizes the different estimated correlation between observables and EE mobility. In Step 3, Assumption 3 ensures that the function $b(X | \gamma)$ is time-invariant, so $\widehat{B}_{i,t} = b(X_{i,t} | \widehat{\gamma}^{R_{i,t}})$ for the entire post-RIP period. Finally, the smaller the share of missing answers in the survey population, the larger the adjustment in Equation (5) needed to guarantee that unobserved heterogeneity has zero mean in the population by Assumption 1. Intuitively, if few suddenly missing answers coincide in time with a huge jump in measured EE, it has to be the case that the records with missing answers had very unusual EE.

A potential concern is that the effect of the RIP may be time-varying, even conditional on respondent group R and on other observables X , violating Assumption 3. Our evidence suggests that this is indeed the case when comparing 2007 and later years, because the behavior of EE_t^m differs. In the imputation regression, we supplement the $\text{RIP}_{i,t}$ dummy with a dummy for the pre-RIP interviews with Blaise software, the middle block (italicized dates) in Table 1, and

¹⁴For the denominator, we restrict attention to records that we can match as described in Section A.A1. The retrospective nature of the EMPSAME 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.

allow the function $b(X_{i,t} | \gamma)$, specifically the parameter vector γ , to differ between the middle and the lower-left (RIP) blocks. So, in Step 2, the regression is run on a constant, observables, two “measurement error” dummies (a Blaise dummy and a RIP dummy) and the interactions of each dummy with a separate flexible function of observables.

C Imputation regressions: specification and results

In Step 2 we specify the function $b(X_{i,t} | \gamma)$ to be linear in the following observables $X_{i,t}$: an aggregate labor market indicator, to be discussed shortly, and dummies for calendar month, Rotation Group (1-2, 2-3, 3-4, 5-6, 6-7, 7-8, denoted by RG1-3 and RG5-7), gender, education (less than HS, HS, Some College, College, Graduate Degree), marital status (Married, Married with Spouse Absent or Separated, Widowed/Divorced, Single), age (16-20, 21-30, 31-40, 41-50, 51-60, 61-70, 71+), major industry (16 major industries, adjusted for breaks to be consistent over time) and major occupation (13 major occupations, adjusted for breaks to be consistent over time). In the Probit regression of Step 1, we use the same covariates except industry and occupation dummies, because estimation of the full specification sometimes fails to converge.¹⁵

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

For this purpose, we choose as our aggregate labor market indicator the observed average EE probability in the same calendar month of the SS records between their first and second rotations (EESSRG1). In order not to restrict the effects of the trend and the cycle in such an indicator to be the same, we first fit a quadratic trend to the monthly series of EESSRG1 over the entire period, and then incorporate the fitted quadratic trend of EESSRG1 and the deviation from it in month t as separate regressors in the vector of observables $X_{i,t}$. We choose SSRG1 as a reference group that is immune, by design, to the effects of the RIP (SS) and of survey attrition (RG1) on the response rate to the EMP-SAME question. After experimenting with many detrending methods, we choose a quadratic trend, and cyclical deviations thereof, until the imputation regression delivers an estimated bias $\hat{B}_{i,t}$ that, once averaged within each respondent group R, shows no residual trend or cyclical variation (see Figure A3 in the Appendix),

¹⁵As mentioned in Footnote 3, the dependent variable equals $DI_{i,t} = 0$ for all members i of RG2 in $t = \text{June 2015}$. For this reason, in that month only, the Probit regression omits RG dummies from the covariates.

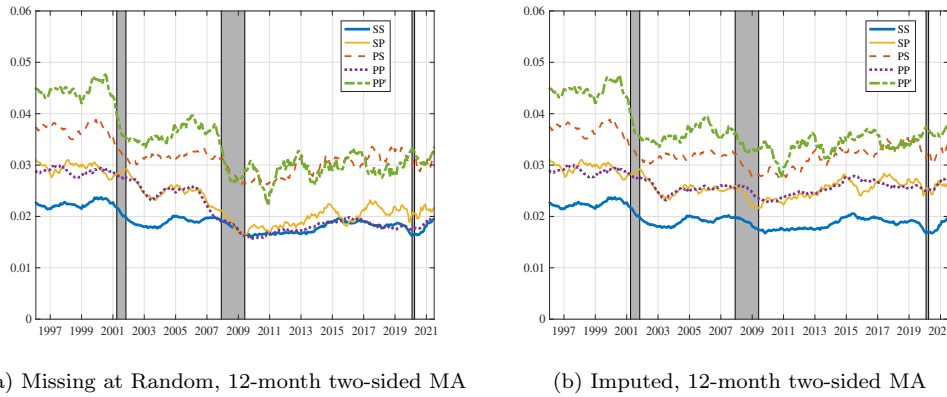


Figure 13. : EE probability by CPS Respondent Group

validating Assumption 3. In this sense, EESSRG1, filtered through the quadratic trend, performs better than other aggregate indicators of labor market conditions that are also immune from the RIP, such as the UE transition probability of the same R group.

We report the results of the Step-2 imputation regression in three tables in the online appendix, and briefly comment on them here. As expected, EE mobility is higher among individuals who are in the first rotation group (thus less selected by job-mobility-related survey attrition), less educated, less attached to spouses, and younger. Mobility is also higher when the average EE mobility of RIP-immune records (EESSRG1) is higher, both in trend and business cycle. $BLAISE_{i,t}$, referring to the records in 2007-2009 not yet exposed to the RIP, signals a drastic level shift down in observed EE mobility, while the impact of $RIP_{i,t}$ works mostly through its interaction with other observables. These two findings indicate that the measurement issue captured by $BLAISE_{i,t}$ is harder to interpret, while the RIP has no impact on the baseline group and operates mostly through selection. We indeed find that the interactions of the two flags, especially $RIP_{i,t}$, with observables $X_{i,t}$, especially age, are often sizable and statistically significant. The declining age profile of EE mobility, which still survives after controlling for many other worker and job characteristics, is much less pronounced after 2007, and even more so after the introduction of the RIP. This finding indicates that the RIP caused a selection out of the valid sample of more job-mobile individuals among young workers, who are more mobile to begin with. That is, $f_{DI}(R, X, \theta)$ is estimated to be submodular in age (which is part of X) and unobserved propensity θ to change employer.

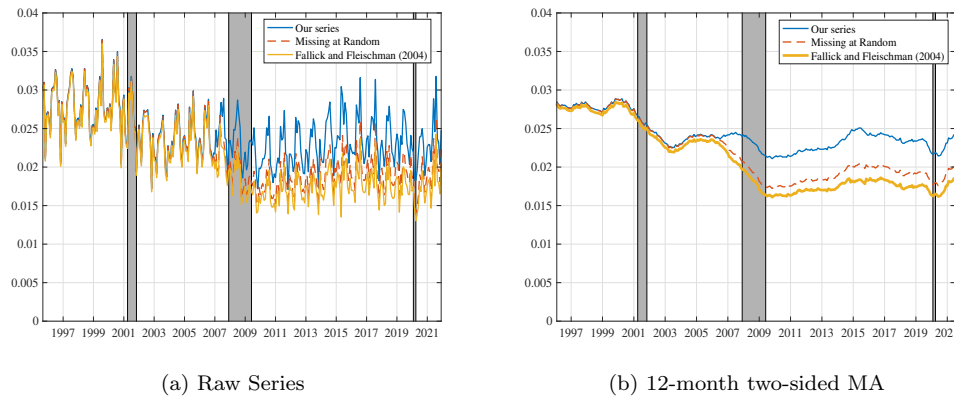


Figure 14. : EE probability: Fallick and Fleischman (2004), Missing at Random, and Imputed

D The imputed aggregate EE probability series

We can finally report our main results. Figure 13 shows, for each respondent group $R \in \{SS, PS, SP, PP, PP'\}$ on which we perform the imputation separately, the time series for the average monthly EE transition probability since 1995, estimated using the Missing at Random assumption (MAR) and our imputation method. To suppress high-frequency noise from sampling error, we plot all monthly time series as two-sided 12-month Moving Averages.

The imputed series, which by construction start diverging from the raw ones after January 2007, are consistently higher, especially for respondent groups SP and PP'. This suggests that respondents who denied permission to share their answers with other household members, thus invalidating Dependent Interviewing questions, including EMPSAME, exhibit observable characteristics that strongly correlate in other records with EE mobility.¹⁶

In Figure 14, we aggregate the group-specific series from Figure 13 and report the main result of our paper, which in part replicates Figure 1a: three time series for the average probability of monthly EE transition in the US since 1995, estimated using the Fallick and Fleischman (2004) method (FF), the Missing at Random assumption (MAR), and our imputation method. In the right panel, all monthly time series are MA-smoothed as before.

By an unfortunate coincidence, measurement issues caused by the January 2007

¹⁶In the online Appendix, we use the delta method to estimate the contribution of the (im)precision in the probability of missing answer imputed from the Probit, $\hat{P}_{i,t}$, to that of the imputed average transition EE probability for each Respondent group. The resulting confidence intervals on our imputed average EE series are extremely tight, except for the PP' group during the pandemic lockdowns of March-April 2020.

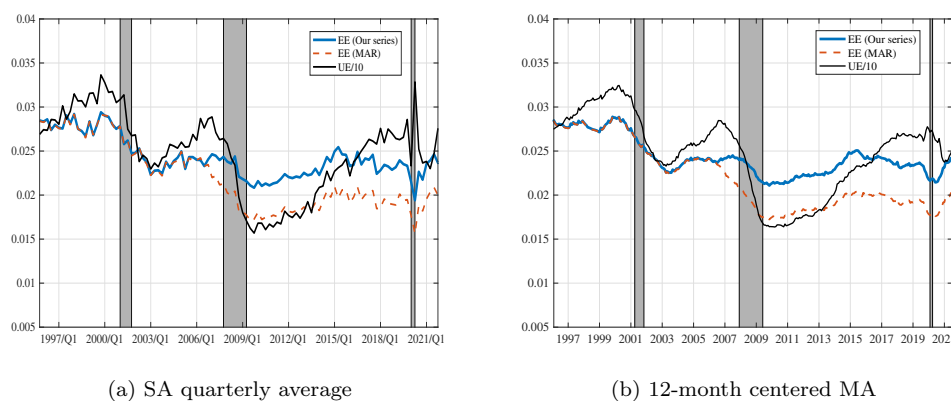


Figure 15. : Comparison of EE Probabilities with UE probability

introduction of the Blaise interviewing software, as revealed by the jump in missing answers to the EMPSAME question in that very month, predate by about a year the onset of the Great Recession. Since the EE transition probability is procyclical, the sharp drop observed around 2007 in the “raw” (estimated according to either the FF or MAR method) EE probability is easily attributed to the recession. Our imputation procedure leads us to conclude that most of the drop was spurious. While the imputed EE probability did fall, importantly, it declined later, and by much less than the raw EE series and the UE transition probability, which declined by about half starting in late 2008, following the financial crisis. The FF/ MAR raw series and our imputed EE series share a weak recovery in 2010-2014, and then a clear rebound, which ends in 2016. Thereafter, our EE series returns to the pre-Great Recession level of about 2.5% and then starts to mildly decline, while the raw series 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 EE turnover appears to have stabilized in the last two decades. As part of the (post-)pandemic “Great Resignation”, our EE measure experienced a sharp spike in 2021, which died out late in the year, when EE returned to pre-pandemic levels. This makes the argument of a declining trend in the 21st century even harder to support statistically.

Note, however, that EE turnover is not an exception in regard to recent trends. In Figure 15, we plot two EE probability series (MAR and our series) along with the transition probability from unemployment to employment (the UE probabil-

ity). One can see that, the UE probability and our EE probability series share very similar trends. In particular, in the post-Great Recession period, both series recovered almost fully to pre-Great Recession levels. The MAR series, on the other hand, diverges from the UE probability over the same period and the gap has been consistently widening. Furthermore, pre-Great Recession the MAR series peaked well before the UE probability (around 2007 in the figure), and fell steadily over the next five years or so. In contrast, our imputed series exhibits roughly a symmetric hump during the period surrounding the Great Recession (between 2004 and 2010) and the UE probability shares a similar symmetric pattern during the same period, although it displays sharper cyclical responses. Importantly, we do not use the UE probability in our imputation at all and there is no reason to believe that the UE probability is also plagued by the measurement issues that affected the measurement of EE transitions. Thus, Figure 15 provides independent evidence that validates our imputation within the CPS.

While the EE rate is procyclical, like UE, it is much less volatile, and tends to stall late in expansions, e.g. 2004-2006 and 2015-2019. As we will see shortly, both facts emerge also from the LEHD-based measure of EE. Moscarini and Postel-Vinay (2023) propose the following interpretation. Each recession slows down the reallocation of workers up the job ladder, and replenishes the stock of mismatched workers at the bottom of it, who are willing to quit. During the subsequent recovery, it takes several years of voluntary quits to deplete this stock. At that point, only robust job postings and frequent offers can fuel further quits. In other words, in the mature phase of those expansions, contacts with open vacancies remained high or even kept increasing, but willingness to change jobs decreased; the supply of potential quits is countercyclical, while the demand is procyclical, and measured quits are the result of these two opposing forces. In light of this interpretation, we can only speculate that the profound sectoral shock associated with the Great Recession (construction, finance) generated a large employment mismatch, that fueled EE even when job openings were scarce. The same has been occurring after the pandemic, an even deeper sectoral shock.

V Comparison with other datasets

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

A Average levels and survey attrition

As mentioned earlier, a limitation of the monthly CPS for our purposes is its address-based nature. If an employed individual moves out of a selected housing

unit to take another job, possibly a household head with the whole household in tow, the survey will lose track of them and miss the EE transition altogether. The same is true of any previously employed individual/household who moves into the same housing unit and thus into the survey, to take another job without any jobless spell. We do not know their employment status before they enter the survey. That is, we need to worry about the correlation between survey attrition and EE mobility.

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

Our first comparison is with the quarterly Longitudinal Employer-Household Dynamics (LEHD), an administrative, matched employer-employee dataset, which contains quarterly reports on total earnings accruing to each US worker from each employer over the entire calendar quarter (Census Bureau (2022*b*)). Unlike the monthly CPS, this source does not suffer from missing answers, but EE transitions still require an imputation, because of a time aggregation bias. Specifically, we know when a worker earned income from two different employers A and B in quarter t , but to label this an EE transition from A to B in quarter t we have to rule out the possibility that there was a jobless spell in between, which the dataset does not report. Hyatt et al. (2014) propose and implement a filter based on a change in quarter t of “main employers,” defined as the main sources of earnings in quarter $t - 1$ and $t + 1$, with at least one of the two also detected at t . This is the methodology adopted by the Census to estimate the LEHD Job-to-Job Flows (J2J) series that we use.¹⁷ Hyatt et al. (2014) report that 15% of measured EE transitions that happen within a quarter, and half of those that happen across adjacent quarters, correspond to a temporary earnings loss of over one month of earnings, compared to the pre- and post- transition quarters, likely due to a jobless spell longer than a month. Bertheau and Vejlin (2022) re-create

¹⁷The J2J rate is currently available in 2000:Q2-2021:Q1 from <https://lehd.ces.census.gov/data/>.

quarterly time aggregation of EE transitions in Danish administrative data, which have exact start and end date of each job and comparable turnover rates to the US. They allow for up to a week of non-employment between jobs, and find that adjacent-quarter transitions are over-estimated by 30%, in a procyclical manner, while within-quarter transitions actually *underestimate* the truth by 14%, a very different outcome than in the US, due to very many transitions that occur in Denmark right at the seam between quarters (March 31, June 30, etc).

To calibrate our monthly CPS estimate of EE to the quarterly level of LEHD, we select CPS individuals who have complete interview histories (from RG1-RG4) and no missing answer to the EMPSAME question at any point in the survey. We focus on cohorts who enter the survey from January 2000 through December 2005, which is the early period covered by the LEHD, and when the RIP or Blaise does not apply yet to the CPS so we can use raw numbers and not our imputed series. In this set, we identify the number of workers who were employed for all four consecutive months RG1-RG4, and estimate their share who experience at least one EE transition during those three pairs of months (quarter). We obtain 5.53%, which is almost identical to the 5.55% average in the quarterly LEHD over the same period.

This congruence is reassuring about the irrelevance of geographical attrition in the CPS for EE measurement. To further corroborate this conclusion, we turn to the SIPP. In principle, unlike the CPS, this representative survey tracks individuals even when they move. In practice, the SIPP also suffers from attrition, but at lower rates than the CPS. We use the 2014 panel (Census Bureau (2014)), when the survey first asked about the reason for a change of address. The 16 possible reasons well illustrate a variety of job-unrelated reasons for the move.¹⁸ Of those who were employed at least part of month t , 1.06% moved within state, 0.21% moved to a different state, and 0.01% moved abroad between months t and $t + 1$. Of the within-state (out of state) movers, 4.54% (respectively, 30.68%) say they moved to take up a new job. Overall, about 0.1% of those who were initially employed changed address to take up a new job. This further suggests that the bias in the CPS due to correlated attrition and EE mobility is quantitatively negligible.

To check whether the SIPP itself well represents the fraction of employed workers who change address, we use the American Community Survey, an annual, large, representative cross-sectional sample of the US resident population. The IPUMS Abacus (IPUMS (n.d.)) tabulates the share of *currently* employed workers who report having moved house in the last year. According to this tabulation, 12.2% moved within state, 2.5% moved from out-of state, and 0.5% moved from

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




Figure 16. : CPS- vs. SIPP-based EE series

abroad per year during 2014-2018 and these shares are nearly constant over this period. These annual numbers correspond almost exactly to the monthly numbers from the SIPP reported above.

The SIPP itself has been used to estimate the average level and time series behavior of EE transitions in the US economy. It turns out that the Census Bureau first applied the RIP to the SIPP in 1998 (Bates, Doyle and Gates, undated; Pascale and Meyer, 2004). To construct EE transitions in the SIPP, we apply the methodology in Moscarini and Postel-Vinay (2016). The EE series starts with the 1996 SIPP panel, after the Survey redesign, and the actual EE transitions can be estimated only starting in April 1996 because of left-censoring, so they are almost entirely impacted by the RIP. The 2014 and 2018 SIPP panels, which cover data from 2014 to the present day, suffer from a radical change in survey design and implementation: due to budget issues, the frequency of interviews declined from thrice to once a year, exacerbating the recall bias of the interviewed, who likely forgot some job switches made by other household members many months before. Indeed, the level of EE transitions in the SIPP drops inexplicably in 2014, both compared with pre-2014 SIPP data and with other datasets, and remains lower. Therefore, we limit our comparison between SIPP and CPS to 1997-2013, which includes the critical 2007-2008 break. We show the results in Figure 16.

Until 2008, the level of EE in the SIPP lies well below the raw (uncorrected, FF) level in the CPS. After 2008, once the RIP is introduced to the CPS too, the gap closes completely, and the SIPP and raw (uncorrected, FF) CPS measures of EE come together. So, it appears that most of the gap between EE rates in the SIPP and the CPS in the 1990s and 2000s, which had been noticed before by other authors, is due to the earlier (by 10 years) introduction of the RIP in the SIPP, with its depressing effect on measured EE. Conversely, our imputed series

remains higher than the SIPP throughout. Note that the SIPP series, just like our imputed CPS series and unlike FF-MAR, shows no trend between early 2000s and 2009-2014.

Unlike in the monthly CPS, it is difficult to discern a sudden drop in measured EE transitions in the SIPP when the RIP was introduced to that survey in 1998. The different structure of the SIPP can explain at least in part this difference in outcomes. First, SIPP interviews occur every four months, staggered, so only 1/4 of the sample is being interviewed (and potentially impacted by the RIP) in each calendar month, and the RIP has bite only at the “seam” between SIPP waves. Second, the SIPP always encodes a job (employer) ID, whether or not Dependent Interviewing applies, while the CPS asks about a *change* of employer and yields a missing observation if the RIP respondent declines. Finally, the definition of EE transition in Moscarini and Postel-Vinay (2016) that we adopt here requires a change not only in job ID, but also in at least one of start date of the job, industry or occupation. This additional filter may have already selected out, before 1998, some SIPP records of the same kind as those directly impacted by RIP after 1998.

B Time-series variation

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

To facilitate comparison with the quarterly, seasonally adjusted LEHD-based J2J series, we take our CPS series, which we seasonally adjust using the Census X13 software, and the seasonally-adjusted JOLTS series, take quarterly averages, and rescale them so that the average level of all series match up for the first three years of the sample.¹⁹ Figure 17 reports the results. The CPS-based FF and MAR

¹⁹We rescale the series because JOLTS quits include those to non-employment, which cannot be separated from those to other jobs, while our CPS series and, more severely, the LEHD J2J series, include both some involuntary but short unemployment spells and direct EE transitions caused by the pre-announced termination of the first job. In order to make our CPS series as comparable with JOLTS quits as possible, in the online Appendix we add to our CPS series the transition rates from Employment to Nonparticipation, excluding those to due to retirement and disability that are tallied separately in JOLTS, and from Employment to Unemployment due to quits (Job Leavers). We find that the resulting augmented CPS-based series significantly exceeds JOLTS quits (reminiscent of Hershbein (2017)’s similar finding for total hires in the CPS exceeding those in JOLTS), and, just like our baseline EE series, shows none of the post-pandemic spike in JOLTS quits, casting doubts on the resulting “Great Resignation” narrative.




Figure 17. : Quarterly EE probability: CPS, LEHD and JOLTS

series start dropping in early 2007, well before the Great Recession, and never recover pre-recession levels. In contrast, our imputed series, as well as LEHD and JOLTS all drop in earnest during the Great Recession, especially in 2008:Q3 when the financial crisis begins, and all recover their pre-recession levels by 2016. We conclude that declining dynamism in US labor markets ends in the 20th century, at least as far as mobility between firms is concerned. In early 2008 our imputed series remains flat, while the other four series are all declining, although not synchronously. The correlation coefficient of each of the three quarterly CPS series in Figure 17 with the LEHD series is 0.6521 (FF), 0.7174 (MAR), 0.8173 (our imputed series). The stronger correlation of our imputed series with the LEHD is close to that shared by all three CPS measures before 2007 (2000:Q2-2006:Q4): 0.8857(FF), 0.8888 (MAR), 0.8910 (our imputed series), providing further evidence in favor of our imputation. The correlation coefficients with the JOLTS quarterly series are 0.4769 (FF), 0.5509 (MAR), 0.7442 (our imputed series).

As further evidence that the drop in the EE probability in the monthly CPS as measured by the FF/MAR method starts prematurely, in 2007, which is precisely when Blaise applies, we can take the peak in EE series to be 2006:Q4 and the trough to be 2009:Q3. By the time the recession began in late 2007, MAR had dropped about half of the total, 5.5% to 4.7%, then from 4.7% to 4% during the recession. In contrast, for JOLTS and LEHD, about four fifths of the peak-to-trough total drop happens during the recession. For example, JOLTS drops in total from 6% to 3.5%, but before the recession that's just 6% to 5.5%, so .5% out of 2.5%, or one fifth.

During the Great Recession, the LEHD and JOLTS series drop proportionally a lot more than our imputed series, and indeed than any CPS series. This




Figure 18. : EE and UE transition rates during the pandemic: Sep. 2019 - Dec. 2021

difference in cyclical response raises the concern that our imputation might be over-correcting the drop due to the RIP. We know, however, that, for the LEHD, the large drop is due, at least to some extent, to time aggregation, described earlier. Short jobless spells, missed by the described procedure to eliminate in the LEHD spurious EE transitions, are more likely in expansion, when then measured EE tend to be exaggerated, and vice versa in recessions. More limited time aggregation exists also in the CPS, because the EMPSAME question does not distinguish between direct EE transitions and very short jobless spells that complete within the month. Regarding JOLTS, quits to nonemployment are likely to be procyclical, because they are less risky at times of high employment, and thus amplify the cyclical volatility of the overall quit rate in the figure. A more detailed analysis on this point is available in the online Appendix.

Moving beyond the critical 2007-2009 period, where changes in Census interviewing overlapped with the Great Recession, we can draw three conclusions from our new evidence, on the long run trend, the cyclical pattern, and the recent behavior of EE mobility during and after the pandemic.

First, the trend. While the time series is too short to draw any rigorous statistical inference on the trend, we can say with some degree of confidence that our imputation eliminates most of the downward trend in EE mobility that the raw data suggest in this century. This conclusion is in line with the behavior of the UE rate as shown in Figure 15.

Second, the cyclical pattern. While the EE rate is procyclical, it tends to stall late in expansions: 1996-2000, 2004-2006 and 2015-2019. This is true also of the LEHD-based measure and of JOLTS quits since 2000 for which these two series are available. In contrast, UE rate and the vacancy/unemployment rate both

kept rising until each cyclical peak. Moscarini and Postel-Vinay (2023) interpret this evidence as follows. After several years of quits, the stock of mismatched workers, who are willing to quit, was depleted, so only robust job postings could fuel further quits. In other words, contacts with open vacancies remained high or kept even increasing, but willingness to move decreased; the supply of potential quits is countercyclical, while the demand is procyclical, and measured quits are the result of these two opposing forces. So we do not take those flat stretches in the EE probability as evidence of an underlying downward trend.

Third, recent events. All current series, once seasonally adjusted and quarterly averaged, drop sharply during the COVID-19 lockdown and the resulting freeze of the US labor market in the second quarter of 2020, and then rebound. As in the Great Recession, the drop is especially pronounced in JOLTS. As part of the (post-)pandemic “Great Resignation”, our EE measure experiences a sharp spike in the summer of 2021, which dies out late in the year, when EE returns to pre-pandemic levels. This makes the argument of a declining trend in this century statistically even harder to support.

In Figure 18 we return to monthly observations and zoom onto 2019-2021. We plot our imputed series and the MAR series, as well as, for reference, the UE transition probability, all seasonally adjusted. The pronounced drop in the EE probability after March 2020 is followed by a strong recovery, which was complete by the summer. EE transitions slowed down again in Fall 2020, in line with the U.S. macroeconomic recovery, although the level of EE in December 2020 was not unusually low. This pattern indicates that the pandemic dramatically delayed EE reallocation during the year (the .6% drop in the EE probability in early 2020 amounts to about one million fewer workers who changed employer per month), but did not significantly change its total volume. While the EE probability declined, in contrast, the UE probability experienced a huge temporary surge due to recalls which do not involve labor reallocation between jobs. In this last recession, EE appears to have been a more meaningful real-time gauge of the pace of reallocation in the US labor market than UE. More generally, this graph illustrates how an accurate and prompt measure of high-frequency EE reallocation can inform policy.

C Comparison with LEHD by demographics

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

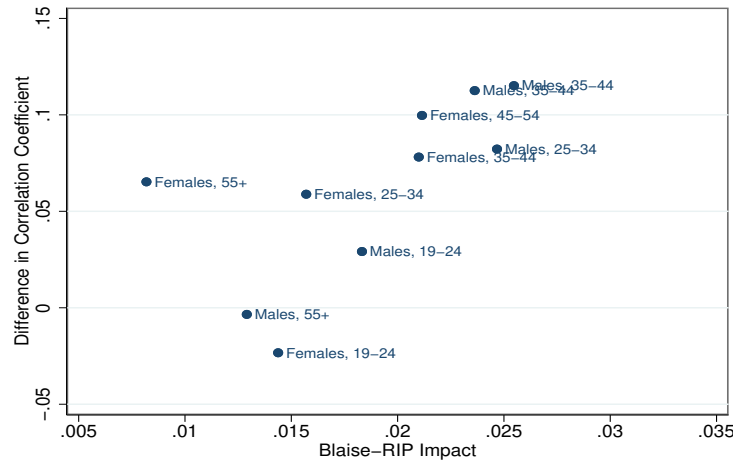


Figure 19. : Improved congruence of CPS series with LEHD due to imputation against impact of Blaise-RIP, by demographic group.

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

Here it is critical to emphasize that, from a conceptual viewpoint, even a single day of unemployment in-between jobs, that the LEHD could never detect, makes a huge difference to earnings and productivity dynamics, if the worker did not know about the new job when they separated from the old one. From a practical viewpoint, whether a very short non-employment spell should be counted or not as an EE transition is a matter of interpretation. Some employed workers line up a new job, and then take some time off before starting it; others lose a job involuntarily, but are lucky enough to find a new one quickly. No dataset allows one to reliably identify this distinction, which is conceptually critical to determine the productivity and earnings implications of the transition. Therefore, researchers have to make some, necessarily dataset-specific, assumptions, which complicate the comparison between datasets with different frequencies.

Besides the difference in time aggregation, the different frequency of observations also complicates a direct comparison between EE in the monthly CPS and in the quarterly LEHD. Multiple monthly EE transitions made by the same individual within a calendar quarter are rare. Therefore, in most cases, each individual contributes at most one transition, monthly and quarterly, to each calendar quarter, and the quarterly average of our monthly CPS-based EE transition probabilities should be directly comparable with the LEHD's measure, except for the different time aggregation. We showed the results of this comparison it at the national level in Figure 17, and now do it by demographics.

The quarterly EE transition probabilities constructed from the LEHD are available also by demographics and some job characteristics from 2000:Q2 to 2020:Q1. Because the monthly CPS has a small sample size, relative to the LEHD, we cannot disaggregate it too finely by demographics, as well as time. As a compromise, we form ten demographic groups, by gender and age (19-24, 25-34, 35-44, 45-54, 55 and up).

First, we estimate the “impact of Blaise-RIP on (the EE transitions of) a demographic group”, as follows. We use all available CPS data, excluding the critical months 2007:M1-2009:M3. We know that through the end of 2006 there was no Blaise-RIP measurement problem, and after March 2009 everybody is subject to the measurement problems, so we want to isolate a clean treatment. For each of the ten demographic groups, we run a separate linear regression of the probability of missing answers (DI=0) on month-of-the-year dummies, to control for seasonality, linear and square terms of time, to capture the trend unrelated to the Blaise-RIP measurement issues, and the dummy that equals zero before January 2007 and one after March 2009. In the online appendix, we show that the regression captures the overall pattern of missing observations well. The estimated coefficients of the Blaise-RIP dummy are placed on the horizontal axis of Figure 19. Note that in this linear probability setting these numbers roughly correspond to the size of the jump in the predicted probabilities between Jan. 2007 and Mar. 2009. There are significant variations across group. Males and middle-age individuals tend to experience a larger “impact” than women and younger people.

Second, using data from the period when LEHD is available, we correlate, across these 10 groups, the quarterly levels of EE rates in the LEHD, seasonally adjusted, with those in CPS (MAR, our series). The correlations range from +0.2 to +0.7 depending on the group and the version of the CPS series. The difference in correlations between our series and MAR is positive for almost all groups and is on the vertical axis of Figure 19.

We can see that, on average, our imputed CPS-based EE series tracks more closely the LEHD EE series than MAR, and this closer congruence is more pronounced the higher the estimated impact of Blaise-RIP, on voiding their answers to the EMPSAME question. We take this evidence as further validation for our imputation, relative to the MAR assumption.²⁰

VI Conclusions

We measure aggregate employer-to-employer (EE) transitions made by workers, without intervening significant jobless spells, in US labor markets. We draw from

²⁰In a previous draft, we ran a separate linear regression of the quarterly EE probability from the LEHD on that from the CPS, across groups and quarters, and a constant, weighting by group employment shares in the LEHD. Plotting the time series of the estimated intercepts and slopes of this sequence of year-by-year linear regressions shows that, relative to the MAR series, our imputed series produces an estimated intercept significantly closer to zero and an estimated slope closer to one at all points in time, establishing a closer congruence between our series and the LEHD by demographics.

the monthly Current Population Survey, the premier resource of information on aggregate labor market dynamics for timeliness and detail. We uncover a drastic increase in the incidence of missing answers to the pertinent survey question (EMPSAME) starting in January 2007, coinciding in time with the roll-out of a new software instrument used to conduct monthly interviews, and predating by about a year the full introduction of new interviewing policy, the Respondent Identification Policy (RIP). We provide evidence that these answers are not missing at random, and these interviewing changes caused a serious permanent downward bias in the standard measure of EE transitions. We propose a model of selection by observable and unobservable worker characteristics, and build on it to impute the missing answers to recover the true aggregate employer-to-employer monthly transition probability. We show that its decline observed during the Great Recession started about a year later and was much less dramatic than the raw, biased series indicates, and had fully recovered by 2015. We conclude that the EE transition rate in the US is procyclical, but less volatile and higher than previously thought, and presents no low-frequency trend in the 21st century. We offer our imputed series as the benchmark measurement of the pace of EE employment reallocation in the US.

Our analysis still faces important limitations. First, we do not know the reason why the new interviewing software caused a deterioration in response rate. More importantly, the share of invalid answers to the EMPSAME question in the CPS was modest but slowly rising even before 2007; at that point, this share experiences a few upward jumps, related to the procedural changes we emphasized, through early 2009, but then continues to rise even 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 the RIP and partially immune to our imputation. In future research, we plan to investigate the causes of these ongoing trends. Getting to the bottom of this measurement issue is especially important in light of the recent debate on declining dynamism in US labor market.

REFERENCES

Adams, David E., James R. Back, John T. Baker II, Candice A. Barnes, Redouane Betrouni, Benjamin C. Bolender, Rosemary Byrne, Stephanie Chan-Yang, Yang Cheng, Dennis Clark, Lisa A. Clement, Jeffrey S. Corteville, Annetta Golding DePompa, Christopher Dick, Khoa Dong, James Farber, Tremika R. Finney, Leslie R. Flores, Brian A. Fouche, Susan S. Gajewski, Susan T. Galeano,

- Aaron J. Gilary, Jennifer Williams Hill, David V. Hornick, John A. Jones, Olga Koly, Chris Kuwik, Mei Li, Kyra M. Linse, Clifford L. Loudermilk, Antoinette Lubich, Tim J. Marshall, Alan K. Peterson, KeTrena Phipps, Jimmie B. Scott, Franklin D. Silberstein, Larry Douglas Sink, Daniel Sommers, Thomas R. Spaulding, Kenisha Staine, Tim Trudell, Katrina L. Wengert, Gregory D. Weyland, Dorinda Allard, Andrew Blank, Thomas Evans, Steven Haugen, Fran Horvath, Karen Kosanovich, Justin McIllece, Stephen M. Miller, Anne Polivka, and Ed Robison. 2019. “Design and Methodology: Current Population Survey - America’s Source for Labor Force Data.” *Census Bureau Technical Paper 77*.
- Bertheau, Antoine, and Rune Vejlin. 2022. “Employer-to-Employer Transitions and Time Aggregation Bias.” *Labour Economics*, 75(5): 102–130.
- Bollinger, Christopher R., Barry T. Hirsch, Charles M. Hokayem, and James P. Ziliak. 2019. “Trouble in the Tails? What We Know about Earnings Nonresponse 30 Years after Lillard, Smith, and Welch.” *Journal of Political Economy*, 127(5): 2143–2185.
- Burdett, Kenneth, and Dale Mortensen. 1998. “Wage Differentials, Employer Size, and Unemployment.” *International Economic Review*, 39(2): 257–273.
- Bureau of Labor Statistics. 2022. “Job Openings and Labor Turnover Survey (JOLTS): Quits rate, Total Nonfarm (Seasonally adjusted) - JTS0000000000000000QUR; Job Openings Level, Total nonfarm (Seasonally adjusted) - JTS0000000000000000JOL.” U.S. Department of Labor, <https://data.bls.gov/cgi-bin/surveymost?jt>, (accessed January 31, 2023).
- Census Bureau. 2014. “Survey of Income and Program Participation 1996-2014.” <https://www.census.gov/programs-surveys/sipp/data/datasets.html>. The Department of Commerce.
- Census Bureau. 2015. “Current Population Survey Interviewing Manual.” CPS-250.
- Census Bureau. 2022a. “Basic Monthly CPS (1994-2022).” U.S. Department of Commerce, <https://www.census.gov/data/datasets/time-series/demo/cps/cps-basic.html>, (accessed on various dates).
- Census Bureau. 2022b. “Longitudinal Employer-Household Dynamics: Version R2022Q4.” U.S. Department of Commerce, <https://lehd.ces.census.gov/data/>, (accessed January 31, 2023).
- Crump, Richard, Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin. 2019. “Unified Approach to Measuring u^* .” *Brookings Papers on Economic Activity*, Spring: 143–238.

- Davis, Steven, and John Haltiwanger.** 2014. "Labor Market Fluidity and Economic Performance." In *Re-Evaluating Labor Market Dynamics. Jackson Hole Economic Policy Symposium Proceedings*, 17–107. Federal Reserve Bank of Kansas City.
- Davis, Steven, and Till Von Wachter.** 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity*, Fall: 1–72.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda.** 2016. "Declining business dynamism: What we know and the way forward." *American Economic Review Papers and Proceedings*, 106(5): 203–207.
- Fallick, Bruce, and Charles Fleischman.** 2004. "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows." Federal Reserve Board Finance and Economics Discussion Series 2004-34.
- Fallick, Bruce, and Charles Fleischman.** 2019. "Data for Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows." <https://www.federalreserve.gov/econresdata/researchdata/feds200434.html>. (Accessed January 31, 2023).
- Feng, Shuaizhang.** 2001. "The Longitudinal Matching of Current Population Surveys: A Proposed Algorithm." *Journal of Economic and Social Measurement*, 27(1-2): 71–91.
- Foster, Lucia, John Haltiwanger, and Chad Syverson.** 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98(1): 394–425.
- Fujita, Shigeru.** 2018. "Declining Labor Turnover and Turbulence." *Journal of Monetary Economics*, 99: 1–19.
- Fujita, Shigeru, and Giuseppe Moscarini.** 2017. "Recall and Unemployment." *American Economic Review*, 107(12): 3875–3916.
- Fujita, Shigeru, Giuseppe Moscarini, and Fabien Postel-Vinay.** n.d.. "Data output from Measuring Employer-to-Employer Reallocation." <https://campuspress.yale.edu/moscarini/data/>.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song.** 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy*, 122(3): 621–660.
- Haltiwanger, John, Henry Hyatt, Lisa Kahn, and Erika McEntarfer.** 2018. "Cyclical Job Ladders by Firm Size and Firm Wage." *American Economic Journal: Macroeconomics*, 10(2): 52–85.

- Hershbein, Brad.** 2017. “The New Hires Quality Index: A Wage Metric for Newly Hired Workers.” W.E. Upjohn Institute for Employment Research Report.
- Hubmer, Joachim.** 2018. “The Job Ladder and Its Implications for Earnings Risk.” *Review of Economic Dynamics*, 29: 172–194.
- Huckfeldt, Christopher.** 2022. “Understanding the Scarring Effect of Recessions.” *American Economic Review*, 112(4): 1273–1310.
- Hyatt, Henry, Erika McEntarfer, Ken Ueda, and Alexandria Zhang.** 2018. “Interstate Migration and Employer-to-Employer Transitions in the U.S.: New Evidence from Administrative Records Data.” *Demography*, 55(6): 2161–2180.
- Hyatt, Henry, Erika McEntarfer, Kevin McKinney, Stephen Tibbets, and Doug Walton.** 2014. “Job-to-Job (J2J) Flows: New Labor Market Statistics from Linked Employer-Employee Data.” *JSM Proceedings 2014*, 231–245.
- IPUMS.** “Abacus.” <https://usa.ipums.org/usa/abacus.shtml>. (Accessed January 31, 2023).
- IPUMS CPS.** 2023. “EMPSAME: Comparability.” (accessed on January 31, 2023).
- Jarosch, Gregor.** 2023. “Searching for Job Security and the Consequences of Job Loss.” *Econometrica*, 91(3): 903–942.
- Lentz, Rasmus, and Dale Mortensen.** 2008. “An Empirical Model of Growth Through Product Innovation.” *Econometrica*, 76(6): 1317–1373.
- Madrian, Brigitte, and Lars John Lefgren.** 2000. “An Approach to Longitudinally Matching Current Population Survey (CPS) Respondents.” *Journal of Economic and Social Measurement*, 26(1): 31–62.
- Mathiowetz, Nancy.** 1992. “Errors in Reports of Occupation.” *Public Opinion Quarterly*, 56(3): 352–355.
- Molloy, Raven, Christopher Smith, Riccardo Trezzi, and Abigail Wozniak.** 2016. “Understanding Declining Fluidity in the U.S. Labor Market.” *Brookings Papers on Economic Activity*, Spring: 183–237.
- Moscarini, Giuseppe, and Fabien Postel-Vinay.** 2016. “Wage Posting and Business Cycles.” *American Economic Review Papers and Proceedings*, 106(5): 208–213.
- Moscarini, Giuseppe, and Fabien Postel-Vinay.** 2018. “On the Job Search and Business Cycles.” Unpublished Manuscript.

- Moscarini, Giuseppe, and Fabien Postel-Vinay.** 2023. “The Job Ladder: Inflation vs. Reallocation.” Unpublished Manuscript.
- Moscarini, Giuseppe, and Kaj Thomsson.** 2007. “Occupational and Job Mobility in the U.S.” *Scandinavian Journal of Economics*, 109(4): 807–836.
- NBER.** 2023. “US Business Cycle Expansions and Contractions.” <https://www.nber.org/research/data/us-business-cycle-expansions-and-contractions>.
- Peracchi, Franco, and Finis Welch.** 1995. “How Representative Are Matched Cross-Sections? Evidence from the Current Population Survey.” *Journal of Econometrics*, 68(1): 153–179.
- Polivka, Anne, and Jennifer Rothgeb.** 1993. “Redesigning the CPS Questionnaire.” *Monthly Labor Review*, 116(9): 10–28.
- Polivka, Anne, Polly Phipps, Christine Rho, and Hugette Sun.** 2009. “The Current Population Survey’s Experience With the Respondent Identification Policy.” *Slides Prepared for May 2009 AAPOR Conference*.
- Postel-Vinay, Fabien, and Jean-Marc Robin.** 2002. “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity.” *Econometrica*, 70(6): 2295–2350.
- Topel, Robert, and Michael Ward.** 1992. “Job Mobility and the Careers of Young Men.” *Quarterly Journal of Economics*, 107(2): 439–479.
- Welch, Finis.** 1993. “Matching the Current Population Surveys.” *Stata Technical Bulletin*, 12(2): 7–11.

APPENDIX

A1 Matching CPS files

Matching monthly CPS files means uniquely identifying records in consecutive survey months that refer to the same individual. In principle, the re-interviewing process in the monthly CPS should allow us to match three-fourths of the sample in any given month to the next month, while one-fourth of the sample exits due to rotation (though individuals in their fourth month can be linked eight months forward). As mentioned, however, various kinds of attrition reduce the fraction of individuals that can actually be matched.²¹

²¹Madrian and Lefgren (2000) and Feng (2001) evaluate in depth the design of the matching criteria of annual (March) CPS records. They build on earlier work in Welch (1993) and Peracchi and Welch (1995).

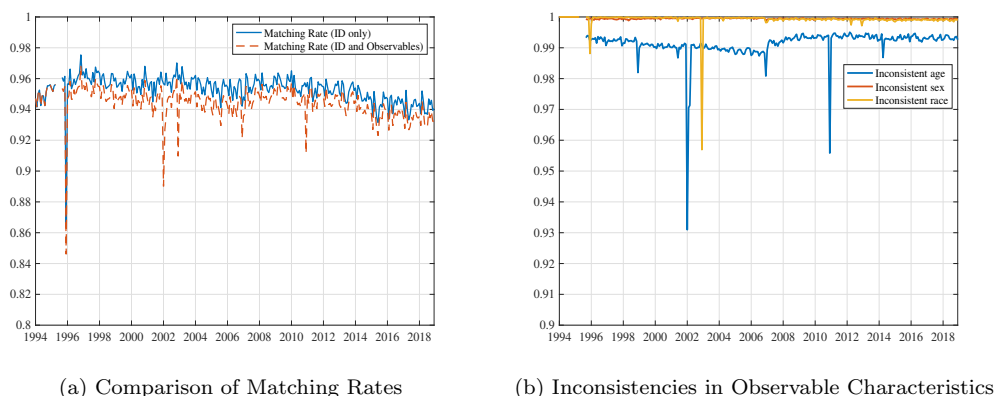


Figure A1. : Matching rates

To match records in January 1994–April 1995, we first take the variable HRHHID, which is 12 digits, and then concatenate it with a 5-digit number, which is in turn created by combining the following three variables: sample number (HRSAMPLE), serial suffix (HRSERSUF) and household number (HUHHNUM). The resulting 17-digit number still does not uniquely identify the household and therefore, even when combined with person line number (PULINENO), the individual. For this reason, following the literature, we also use the individual’s age, gender and race to establish an individual match.²²

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, HRHHID is 15 digits, and its three additional digits, along with the 5-digit number formed by HRSAMPLE, HRSERSUF, and HUHHNUM as before,²³ generate a 20-digit number that uniquely identifies the household. Individuals within the household can then be identified by PULINENO without using observable characteristics. In fact, after September 1995 these observable individual characteristics are likely to generate “spurious mismatches,” because the Census Bureau occasionally “scrambles” respondents’ age information, and more generally because these characteristics may be measured with error. ID variables are arguably more fundamental to the entire survey and thus mistakes in coding the ID variables are likely to be rare or to be eventually corrected before the data is made public.

Figure A1(a) presents the probability that a respondent who appears in the month- t micro data in Rotation Groups 1-3 or 5-7 also appears in the month $t + 1$

²²We allow for age to increase by one year between the two months.

²³Starting in May 2004, this five-digit part, named HRHHID2, is directly available from the data.





Figure A2. : EE probability and quadratic trend: Self-Self between the first and second months of the survey


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 A1(a) 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 A1(b), we present (the complementary) 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.

A2 Imputation Regression: Validation


To validate the key Assumption 3 for our imputation, Figure A2 shows the fit of a quadratic trend to our aggregate labor market indicator, the observed average EE probability in the same calendar month of the SS records who are in the first (and second) rotation (EESSRG1). The quadratic trend and the deviations from it enter separately the imputation regression. Then Figure A3 illustrates the average estimated bias for each month and each respondent group. By construction, the bias is zero before 2007. It is clear that 2007 is different from



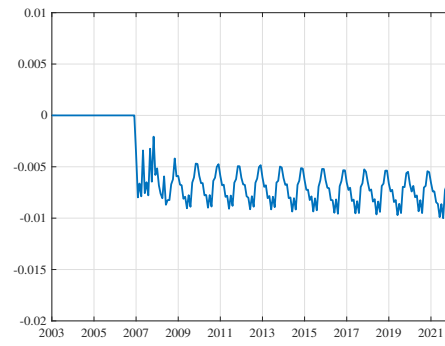
(a) SS




(b) SP



(c) PS



(d) PP



(e) PP'

Figure A3. : Estimated average bias $\mathbb{E}_i[\widehat{B}_{i,t}]$

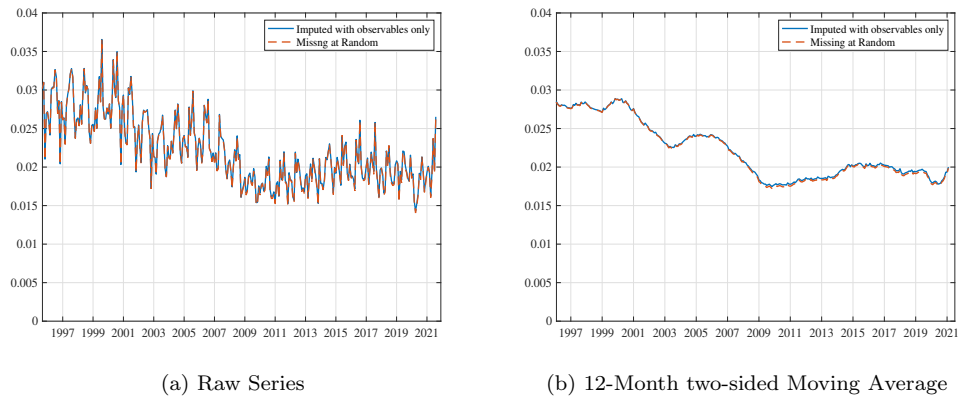


Figure A4. : EE probability: Missing at Random vs. Imputed by observables only

later years, when the average bias settles into a very regular seasonal pattern, with no visible residual trend and cycle, except for a very small decline for the PP group and a slightly hump-shaped pattern for the PP' group. The seasonal pattern of the bias indicates that the underlying seasonal pattern of EE, clearly visible in pre-2007 data, changed permanently with the introduction of the RIP in 2008, and settled into a different, but equally regular, pattern thereafter, further evidence of selection by unobservables. Over the entire post-2007 period, the bias averages approximately zero for the SS group, and is otherwise negative, reflecting the reduction in measured EE due to the correlation between EE mobility and non-response rate. The bias grows in size moving from the PS to the SP group, then further for PP and is largest for PP'.

A3 Imputing missing records by observables only

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

In Figure A4, we compare this series with the one based on the MAR assump-

tion. 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.