A Longitudinal Analysis of the Current Population Survey: Assessing the Cyclical Bias of Geographic Mobility

Christopher J. Nekarda
University of California, San Diego

First version: 12 September 2006
Current version: 27 November 2008

Abstract

This paper assesses the implications of geographic mobility for the measurement of U.S. labor market dynamics using the Current Population Survey (CPS). Because the CPS does not follow individuals that move, estimates may be biased if the labor market behavior of movers differs systematically from that of nonmovers. I create a new database, the Longitudinal Population Database (LPD), that utilizes all longitudinal information in the CPS to form a panel data set. I use the LPD to identify persons who move and therewith estimate a bound on the bias from geographic mobility. I find that the cyclical bias arising from geographic mobility is small. At business cycle frequencies, the difference between the separation hazard rate calculated from the entire CPS sample and from a subset that are known not to have moved never exceeds 4 percent. There is little effect of mobility on the job finding hazard rate. I conclude that geographic mobility does not significantly affect CPS labor market dynamics.

JEL codes: J61, E32, E24, C22
Keywords: geographic mobility, Current Population Survey, separation, job finding, hazard rate, bias

1 Introduction

Many interesting questions about the U.S. labor market are longitudinal in nature. That is, they require observations for the same individual or set of individuals at different points in time. Examples of such research are the dynamics of gross flow of workers, occupational and job mobility, the behavior of real wages over the business cycle, and the decision to migrate.

Economists generally view geographic mobility as a means of reallocating resources, in this case labor, to more efficient uses.[1] Typically 70 percent or more of people who move indicate having moved for economic reasons and up to 50 percent of those moves occurred because of a job separation.[2] In particular, researchers find a positive relationship between unemployment and geographic mobility, consistent with labor reallocation.[3]

The link between labor market dynamics and mobility has important economic and public policy consequences.[4] It also has important implications for the *measurement* of labor market dynamics, particularly when using the Current Population Survey (CPS).[5] Specifically, the CPS does not follow individuals that move away from a sample address, possibly creating a bias in longitudinal measurements. Because of the strong relationship between unemployment, job separation, and mobility, there is concern that the dynamics captured by the CPS may be biased from sample attrition related to geographic mobility.

A proper assessment of this concern requires a new approach to longitudinal research using CPS data. Although the CPS is typically used in cross section, such as when calculating the unemployment rate, an individual's responses in the CPS can be matched longitudinally. Two common uses of this longitudinal feature are to match individuals from one month to the next and to match individuals from one year to the next in the CPS Annual Demographic Supplement. The longitudinal continuity allows researchers to observe changes in individuals' labor force status, income, hours worked, and many other characteristics.

Although these applications exploit the CPS's longitudinal capabilities, they do not make full use of the longitudinal information available. I create a new database that captures all possible longitudinal information in the CPS. Rather than organize the data by month, as the CPS does, I define a *person* as the

1. See Greenwood (1975) for a survey on mobility.
2. Lansing and Morgan (1967); Bartel (1979).
3. Bartel (1979); Schlottmann and Herzog Jr. (1981, 1984).
4. See, for example, Bartel (1979); Topel and Ward (1992); Kletzer (1998); Farber (1999); Gottschalk and Moffitt (1999); Holzer and LaLonde (1999); Neal (1999); Moscarini and Thomsson (2008).
5. Katz et al. (1984); Dahmann (1986); Welch (1993); Fitzgerald et al. (1998); Gottschalk and Moffitt (1999); Neumark and Kawaguchi (2004).

fundamental unit. For each person, I combine all CPS interviews to form a mini-panel containing the largest collection of monthly observations that could possibly have come from the same person.

This database, the Longitudinal Population Database (LPD), contains the complete interview history for every person surveyed by the CPS over 1976–2006. The LPD contains data for over 10 million individuals who together are representative of the U.S. civilian noninstitutional population. Over 65 percent of these individuals have a interview history of at least four continuous months and about 4.5 million persons have a complete history of 8 observations.

The LPD also provides excellent information on mobility. About 20 percent of addresses in the LPD have at least one change in household. Because the LPD contains the entire history of each address in the sample, it is possible to distinguish between individuals that move ("movers") and those that do not ("stayers"). Also, since many movers spend at least four months in the sample, the LPD records their demographic characteristics and a meaningful history of labor force behavior. Furthermore, because the selection of an address for sampling is independent of the decision to move, the LPD provides a true random sample of movers. This allows a meaningful comparison of demographic and labor force characteristics of individuals who move against those who do not.

I use the LPD to assess whether geographic mobility biases labor market dynamics measured by the CPS. Comparing the populations of movers and stayers reveals minor differences in the sex, race, and education of movers but finds large differences in age and marital status of movers compared with stayers. This confirms a well-known feature of geographic mobility, the age selectivity of migration, which identifies a decline in mobility with age.[6]

Also consistent with earlier research, there are important differences in the labor force status between movers and stayers. I find that the proportion of unemployed persons in the population of movers is 60 percent greater than that in the population of stayers. In addition, separations to and accessions from unemployment are twice as frequent among movers compared with stayers. Expressed as separation and job finding hazard rates, movers and stayers to do not differ significantly in job finding rate but movers have a considerably higher separation hazard rate than stayers.

To assess the cyclical effect of geographic mobility, I construct a counterfactual CPS series using only the population of stayers—that is, assuming no mobility. Comparing this counterfactual series with the actual series estimated from the entire population provides a bound on the bias from geographic mo-

6. Gallaway (1969); Schlottmann and Herzog Jr. (1984); Tucker and Urton (1987); Peracchi and Welch (1994).

bility. The bias in the separation hazard rate moves countercyclically, implying that the separation hazard rate calculated using the entire CPS sample will appear too *a*cyclical. However, the magnitude of the bias at business cycle frequencies—the difference between the cyclical component of separation hazard rates in population of stayers and in the entire sample—never exceeds 4 percent. There is little effect from geographic mobility on the job finding hazard rate.

The small cyclical bias can be reconciled with the substantial difference in separation hazard rates between movers and stayers by recognizing the distinction between out-movers and in-movers. The logic for a bias arising from geographic mobility bias is based on sample attrition: individuals that leave the sample are not followed. But there are equally as many people who move *into* the CPS sample as leave it and the differences between the two types of movers are small relative to stayers. Thus, the cyclical bias from geographic mobility is small because people who move out of one address tend to be replaced by similar people elsewhere in the country.

The paper proceeds as follows. Section 2 briefly describes the Current Population Survey, highlighting aspects important for longitudinal matching, and explains the fundamental units of longitudinal analysis. Section 3 uses the LPD to assess the potential bias in the CPS due to geographic mobility. Section 4 explores the robustness of the bias exercise. The final section concludes.

2 The Longitudinal Population Database

The Current Population Survey (CPS) traces its conceptual origins back to the 1930s, when the first monthly national survey to directly measure unemployment began. The modern CPS began in 1948 as the continuation of that survey. The CPS is a monthly survey of about 50,000 U.S. households conducted to gather information about the domestic labor force. Sample households are selected at random and surveyed 8 times over sixteen months. The household rotation design was implemented to maximize continuity from month to month and year to year and to decrease the variance of survey estimates. An additional benefit of the design is that the CPS contains a wealth of longitudinal information.

Starting in the 1980s, the Census Bureau began publishing public-use microdata files containing the outcome of every CPS interview. With this information, researchers started using the CPS to explore longitudinal questions. The publicly-available CPS data are not, however, readily usable for comprehensive longitudinal research. The goal in creating the LPD is to capture all possible longitudinal information on an individual from the underlying monthly CPS surveys. The CPS is a repeated cross-section, organized by month; the LPD uses the person as the fundamental organizing unit. The LPD turns the CPS data into a panel—that is, records the complete interview history of every person surveyed. Although there is a relatively large literature about matching CPS records, previous discussions have focused on month-over-month matching.[7]

A common concern in longitudinal research using CPS data is the large number of unmatched records.[8] Roughly 30 percent of observations cannot be matched from one month to the next. Most nonmatches result from the CPS's rotating sample design, which allows at most 75 percent of individuals to match across successive months. Of observations with the potential to match, roughly 6 percent do not match—over 10 million persons a month. Viewing these missing observations in the context of their complete interview history allows missing observations to be more easily classified.

Despite these shortcomings, the CPS is an excellent survey for economic research because it is a large, random sample from the U.S. population and is the most representative sample available at this frequency. Other databases from surveys, such as the Longitudinal Research Database (LRD), the Business Employment Dynamics (BED), and the Panel Study on Income Dynamics (PSID), contain longitudinal information on their populations.[9] The LRD and BED are not ideal for studying U.S. labor market dynamics because both are surveys about jobs at production establishments and not about individuals—the same person can be employed in more than one job and those not employed are not represented. In addition, they are conducted only annually or, at best, quarterly. The PSID is more appropriate for labor force research, however it is a substantially smaller sample than the CPS and is conducted only annually.

Another survey, the Job Openings and Labor Turnover Survey (JOLTS), provides monthly data on flow of hires and separations at U.S. firms. It began in December 2000 and thus provides a relative short period compared to the CPS. In addition, there are several well-documented discrepancies between aggregate estimates from JOLTS and those from other data sources.[10] In particular, the magnitude of hires and separations in JOLTS are surprisingly small com-

---

7. See Katz et al. (1984); Abowd and Zellner (1985); Hogue (1985); Hogue and Flaim (1986); Poterba and Summers (1984, 1986); Chua and Fuller (1987); Welch (1993); Peracchi and Welch (1994); Hausman et al. (1998); Madrian and Lefgren (2000); Shimer (2007); Moscarini and Thomsson (2008). Feng (2001) also evaluates matches using the complete interview history, but only matches the 1998 and 1999 CPS March Annual Demographic Supplement.

8. Abowd and Zellner (1985); Hogue (1985); Hogue and Flaim (1986); Welch (1993); Feng (2001); Moscarini and Thomsson (2008).

9. Dahmann (1986) discusses using panel data to study geographic mobility.

10. Faberman (2005).

pared to similar measures in other data sources.[11] More importantly, however, the JOLTS is a survey of firms, not workers. It does not include demographic information and is not suitable for studying geographic mobility.

A final U.S. household survey, the Survey of Income and Program Participation (SIPP), is also suitable for studying labor market dynamics and mobility. The SIPP is an ongoing longitudinal survey designed to study longer-term effects of income and government program participation. The SIPP panels last for between one and four years, a substantial improvement over the CPS, however the sample sizes are considerably smaller. Additionally, it is difficult to construct aggregate time-series estimates from the SIPP.[12] However, the SIPP follows individuals that move away from the initial survey address, making it ideal for studying mobility.[13]

2.1 Constructing the LPD

Administered jointly by the U.S. Census Bureau and the Bureau of Labor Statistics (BLS), the CPS surveys 50,000–60,000 households every month from all 50 states and the District of Columbia. It collects complete demographic and labor force data on all persons aged fifteen or older, but records basic information for all household members. Persons on active duty in the U.S. Armed Forces and persons in institutions are not eligible for survey.

The Census Bureau publishes microdata files containing the outcome of every CPS interview beginning with January 1976. The microdata undergo a complex editing and reorganization process to ensure longitudinal continuity and then are combined to create the LPD. This section briefly describes how the LPD is constructed. Appendix A provides a detailed description and technical information.

Despite common use of the word "household," the CPS is, in fact, a survey of *addresses*. The CPS is a multistage stratified sample of addresses from 792 sample areas in the United States.[14] Housing units are sampled from address lists generated from the Decennial Census of Population and Housing and updated for housing built after the census. The sample is drawn once per decade using information from the most recent Decennial Census.

The CPS uses a rotating sample to minimize variance, both between months and between households, as well as reduce the burden on respondents. Each

11. Recent work by Davis et al. (2008) devises a correction for the JOLTS data.

12. Nekarda (2008).

13. Neumark and Kawaguchi (2004) use the SIPP to study of how directly adjusting for geographic mobility compares to the typical Heckman (1979) selection correction.

14. Bureau of Labor Statistics (2002), chapter 3.

address selected for the sample is surveyed for four consecutive months, not surveyed for the next eight months, and then surveyed again for the next four months. It then leaves the sample permanently.

An address is identified over time by its month in sample (MIS) designation, which corresponds to the number of times the address is scheduled to be surveyed. Figure 1 shows the relationship between the MIS and the calendar month of the survey and also the sample rotation. The 4-8-4 rotation pattern enables up to 75 percent of units to match from one month to the next and 50 percent to match from year to year. The large continuity between households across time permits sophisticated longitudinal analysis using CPS data.

The first time an address enters the sample it is visited in person by a Census Bureau field representative to establish whether it is eligible for survey. To be considered eligible the housing unit at the address must be occupied by at least one person eligible for interview (a civilian who is at least fifteen years old and does not usually reside elsewhere). At eligible housing units, the surveyor initiates the CPS interview.

Ineligible addresses are recorded as a noninterview. A *type C noninterview* occurs if the address is permanently ineligible for interview. This condition arises if the housing unit has been converted to a permanent business, condemned, or demolished or if the address falls outside the area for which it was selected. The address is never visited again. A *type B noninterview* occurs if the address is intended for occupancy but is not occupied by any eligible person. Such units are typically vacant, but also include those occupied entirely by individuals not eligible for interview. Type B addresses may become eligible in the future and are thus visited for all eight months that the address is in the sample.

The previous two types of noninterview occur when no one from the civilian noninstitutional population resides at the selected address. Such locations are not considered part of the CPS sample. The third type, a *type A noninterview*, occurs if the address is eligible for a CPS interview but no useable data are collected. This can arise because the occupants are absent or otherwise unavailable during the interviewing period or refuse to participate in the interview. These noninterviews are considered part of the CPS sample. However, because no information about the current occupants is collected, the sample weight of similar nearby units is increased to compensate. The type A condition is considered temporary and the address is visited in all succeeding months.

The BLS assigns each household a scrambled identifier to ensure confidentiality but still permit longitudinal matching. For data after 1994, when the CPS was substantially redesigned, the household identifier is globally unique. Prior to 1994, however, it is only unique across two months for households in the same rotation group. I develop an algorithm to identify households and generate a globally-unique household identifier.[15]

In addition, the BLS periodically changes the scrambled identifier for households. This is disruptive for longitudinal matching in the LPD. For simple month-over-month matching, a change of household identifier prohibits a match only for the month in which the change occurred; all preceding and subsequent months match. However, because the LPD matches an individual across sixteen months, an identifier change disrupts longitudinal continuity for the entire history. Authors either report a missing value for the month where matching was impossible or construct a moving-average across months that do match.

A second challenge in constructing the LPD is ensuring longitudinal consistency. Over the thirty-year period for which microdata are available, the data definitions change 17 times. I develop a consistent set of definitions for categorical variables (e. g., race, educational attainment, or occupation) for the entire LPD.

After creating longitudinally-consistent variables and unique household identifiers for every month, the data are combined together to form the LPD. The LPD has over 53 million observations covering the period 1976–2007, or approximately 140,000 observations per month. The smallest month has just under 97,000 observations and the biggest month almost 160,000.

2.2 Longitudinal Units of the LPD

The objective of the LPD is to construct a complete longitudinal record for every person in the CPS. The CPS, however, is a probability sample of addresses, not individuals. Therefore, constructing a person's longitudinal history begins with the interview history at the address level. In any month, an address is occupied by a single household. But households can move into and out of an address during its time in the sample, generating a difference between the household and the address. Each household consists of one or more individuals. As with addresses, individuals may move into and out of a household. Thus each individual must be identified longitudinally in relation to her household and address. Figure 2 shows the hierarchical relationship between addresses, households, and persons.

An *interview history* is the collection of all monthly observations from a particular unit (address, household, or person). The address is the basic unit. All households and persons from an address inherit the same address interview history. The household is subset of the address. All individuals within

15. Feng (2001) develops a similar procedure to exploit the pattern of sample rotation.

a household share the same household interview history but each household has a unique interview history. The finest unit is the person. Each person has a unique interview history. Table 1 provides example interview histories for different longitudinal situations encountered in the CPS. This table will be referenced throughout the following subsections.

2.2.1 Addresses

Each sample address is scheduled for 8 interviews by a Census Bureau field representative. An *address observation unit* (AOU) is the collection of interviews conducted at an address during its time in the CPS sample. An AOU can have at most 8 observations, but addresses found permanently ineligible (type C noninterviews) will have fewer than 8. Many type C noninterviews are determined on the first interview or following a type B noninterview. Example 3 in table 1 shows the interview history for an address with a type C noninterview. There are about 3.7 million unique AOUs in the LPD (table 2).

2.2.2 Households

Because an AOU spans sixteen months, including eight months without being surveyed, it is possible for more than one household to occupy the address during its time in the CPS. Households that move during the survey are not followed by the CPS; instead the replacement household, if any, is surveyed for the rest of that address's time in the sample. A *household observation unit* (HOU) is the largest collection of observations within an AOU that can possibly come from the same household.

Because individuals are identified within their household, AOUs must first be examined to identify unique households. In most cases an AOU contains only one household, but some AOUs have at least one change of household. A household change can occur in 4 ways:

H1. The original occupants of the address move out and a replacement household moves in with no intervening vacancy recorded.

H2. The original occupants of the address move out and a replacement household moves in but with an intervening vacancy.

H3. The original occupants of the address move out and are not replaced during the address's tenure in the CPS sample.

H4. The address is initially vacant but a household moves in before the address has rotated out of the sample.

The household change in case H1 is straightforward. The replacement household is identified as a household by the CPS. There is no noninterview recorded in the AOU, however it must be partitioned into two HOUs to reflect the change in household. Individuals associated with the original HOU are replaced by the new occupants. Example 2 in table 1 demonstrates such a situation.

In case H2 the replacement household is often not identified as part of a new household. Accordingly, the LPD creates a separate HOU any time a string of completed interviews within an AOU is interrupted by one or more type B noninterviews. Example 5 in table 1 depicts such a situation. The AOU contains a type B noninterview at MIS 4. The first HOU within this AOU contains the first three observations; the remaining completed interviews are assigned to the second HOU.

When a previously-occupied housing unit is found ineligible during all remaining months in the CPS (case H3), the subsequent type B noninterviews are discarded. This is depicted in table 1, example 6. Similarly, when an address is found initially ineligible but subsequently interviewed (case H4) the initial type B observations are discarded. Line 7 of table 1 shows an example of case H4 and the resulting HOUs. Both cases, however, identify households that have moved.

Over 80 percent of AOUs in the LPD have no household change (table 2). These households are known not to have moved during their tenure in the CPS. This does *not*, however, imply that these HOUs have no noninterviews. Type A noninterviews are permitted within an HOU and do not imply mobility. The remaining addresses, just over 19 percent, record a change of household during their tenure in the CPS.

Household changes interrupt longitudinal continuity of the observation unit. For research where continuity is important, such as calculating gross labor force flows, these interruptions reduce the number of observable transitions. For other avenues of research, however, these household changes are beneficial. In particular, a change within an AOU identifies a household that has moved.

On average, each address is occupied by 1.14 households over its sixteen months in the CPS sample. The implied rate of annual mobility, the probability that a household does not reside at the same address one year later, is 14.7 percent.[16] This rate is consistent with the annual rate of geographic mobility estimated by the Census Bureau using the CPS Annual Demographic Supple-

16. The LPD contains 4,160,835 unique households at 3,646,370 unique addresses, yielding 1.1380 households per address (table 2). This implies an annual rate of mobility equal to $1 - (1.1380/16 \times 12) = 0.1465$.

ment. U.S. Census Bureau (2007) reports the average annual mobility rate over 1976–2007 is 14.9 percent.

2.2.3 Persons

Each household has one or more persons residing there. A *person observation unit* (POU) is the largest collection of observations within an HOU that can possibly come from the same person. Because the POU is a subset of the HOU, all POUs within that HOU also terminate when an HOU ends. Example 2 in table 1 demonstrates this: the POU for the person in household 1 terminates when the second household begins.

Also, because individuals can move into and out of a household, each POU can have a different interview history from its associated HOU. Consider, for example, a college student living with her parents during summer: she is counted in the household for interviews conducted during the summer, but her POU terminates when she returns to school. Such a case is shown for the 2nd person in example 8 (table 1).

There are 10.6 million unique POUs in the LPD. The CPS collects full demographic and labor force information only for persons over fifteen years old. For those younger than fifteen, only information on sex, race, and age is collected. There are 2.3 million POUs for persons aged fifteen years and younger. These POUs are not included when studying mobility.

2.3 Longitudinal Statistics from the LPD

How useful the LPD is depends on how much meaningful longitudinal information is contained within the POUs. This section provides a detailed analysis of the POUs and reveals the large amount of longitudinal information contained in the LPD.

For each POU I calculate the number of attempted interviews and the number of completed interviews. For example, the individuals in example 1 from table 1 both have 8 attempted interviews and 8 completed interviews. In example 2, person 1 from household 1 has 4 attempted interviews and 4 completed interviews. The individual in example 3 has 5 attempted interviews, all completed.

Table 3 reports tabulations of the number of attempted and completed interviews for all POUs. POUs are weighted by the average CPS sampling weight for the POU.[17] The column totals (bottom) are the share of POUs with that number of completed interviews. The row totals (right) are the share of POUs

17. See section 2.4 for details.

with that number of attempted interviews. Thus cells on the diagonal are POUs with no noninterviews; these contain the most longitudinal information possible. The sum of the diagonal elements, the share of POUs without missing observations, makes up 94 percent of the LPD.

Because a POU combines two blocks of consecutive monthly interviews, it is also important to identify the number of consecutive months of longitudinal information. The bottom right cell shows POUs with 8 completed interviews, that is, two four-month blocks. It is also the single largest cell, accounting for 31 percent of all POUs.

But many more POUs have at least one block of four months. The next largest cell in table 3 is for 4 completed interviews out of 4 interviews, comprising 26 percent of POUs. Persons with a block of four interviews are important for studying mobility, because often the other block is missing because of a move. The LPD contains about 5.8 million POUs, just under 60 percent of all POUs, with either 4 or 8 completed interviews and no noninterviews.

2.4 Match Validity

The standard procedure in the literature is to match observations from one month to the next using household and person identification variables and then validate these matches using supplementary demographic characteristics.[18] A failure of any criterion invalidates the match.

The LPD allows for much more sophisticated evaluation of matched observations. Instead of evaluating the match just from one month to the next, the entire the interview history can be used. I develop a measure that evaluates each month against all other months for a person, rather than simply month-over-month.

For example, consider a man who is mistakenly classified as a woman for one month of his tenure in the CPS. The standard validation procedure would potentially discard 2 matches (1/3 of the total possible) from this simple mistake (one match on either side of the classification error). One failed match criterion over 8 observations on a person, is very likely to be a clerical mistake and not an invalid match. My method evaluates each month using *all* longitudinal information for the person. In particular, responses for each month is evaluated against those in all other months.

I evaluate a match's validity according to 3 criteria

1. Sex: a person's sex should not change over the POU.

18. For example, Madrian and Lefgren (2000) consider sex, race, age, and educational attainment. Shimer (2007) and Moscarini and Thomsson (2008) use sex, race, and age.

2. Race: a person's race should not change over the POU.
3. Age: a person's age should not change by more than 2 years over the POU.

To formalize, let s_{it} indicate the sex recorded for person i in month t. Similarly, let r_{it} and a_{it} be the recorded race and age in month t. Person i has T_i valid observations in the LPD. The validity score V of the month t observation for person i is

$$V_{it} = \frac{1}{3T_i}\sum_{j=1}^{T_i} I(s_{it} = s_{ij}) + I(r_{it} = r_{ij}) + I(|a_{it} - a_{ij}| \leq 2), \tag{1}$$

where $I(\cdot)$ an indicator function that is 1 if the statement is true and 0 otherwise. For a person with only one observation, $V_{it} = 1$.

If all criteria match for all observations, $V_{it} = 1$ for all t. In the example above, each month's score falls because of the failure of the sex criterion. However in month where sex was female, V is lower still, because $I(s_{it} = s_{ij}) = 0$ for all other months. Thus, this method penalizes all of a person's observations for a single failure; the month with the discrepancy is penalized more.

I treat V_{it} as representing the probability of valid match and adjust the person's month t sampling weight, ω_{it}, by that probability to get the validity-weighted sampling weight $\nu_{it} = \omega_{it}V_{it}$. All population estimates are calculated using this adjusted sampling weight. Thus, each labor force transition is effectively weighted by the "probability" that it came from the same person.[19]

The average validity score in the LPD is 0.9604 when taken over all observations and 0.9930 when taken over nonmissing observations (those with positive CPS sampling weight). This confirms that most matches directly identified by the CPS are valid. In addition, since only the latter group enter population totals, the observed match quality is very high. The results that follow are robust to using other match validation procedures; see section 4.1.

3 Geographic Mobility

Geographic mobility has important implications for the *measurement* of labor market dynamics, particularly when using the CPS. Specifically, the CPS does not follow individuals that move away from a sample address, possibly creating a bias in longitudinal measurements. Because of the strong relationship

19. Feng (2001) evaluates the probability of a valid match conditional on sex, race, age, and marital status using Bayes' rule. This still, however, leads to a binary accept-reject decision.

between unemployment, job separation, and mobility, there is concern that the dynamics captured by the CPS may be biased from sample attrition related to geographic mobility.

The argument that geographic mobility can bias longitudinal measurements is usually phrased in terms of sample attrition: some event, possibly related to the business cycle, causes a household to move out of the CPS sample. Therefore, because the CPS does not follow those individuals that leave the sample, there may be a cyclical bias from geographic mobility.

Sample attrition is not, however, the only type of mobility observed. As section 2.2 emphasized, a change of household at an address can occur 4 different ways, only 1 of which (H3) is pure sample attrition. In fact, the LPD identifies roughly equal numbers of persons moving into and out of the sample. Thus, the language of "sample attrition" is not the correct way to describe geographic mobility in the CPS. Instead, I describe mobility in terms of "out-movers" and "in-movers". An *out-mover* is a person who permanently leaves an address during its tenure in the CPS sample. An *in-mover* is a person not originally present who joins at an address during its tenure in the sample.

3.1 Identifying Geographic Mobility

Geographic mobility is identified using the interview histories in the LPD. Using the full longitudinal history of a person allows me to identify persons that move separately from persons with missing observations arising from some other reason. Two types of mobility can be identified. The first is a complete change of household. This is the most common, accounting for 70 percent of movers.

The population of movers is identified as the set of observations for which the interview history of the HOU differs from that of the AOU. This definition captures all mobility events described by cases H1–H4 and combinations thereof. Mobility is not identified simply based on the number of observations in an HOU nor the existence of missing observations. Instead, mobility is identified using the LPD by the relationship between the HOU and its AOU. For example, line 9 of table 1 shows a case where no interview was recorded in MIS 7. However because the interview history for the AOU and HOU are identical, this household is considered a nonmover. All AOUs with at least one valid observation that have a type B noninterview are identified as movers.

In addition to households that move, individuals can move into and out of households. Examples 8 and 9 in table 1 show such cases. Individuals that move into and out of an HOU are *not* included in the population of stayers. Individual mobility—that is, not associated with a household change—accounts for 23 percent of movers. The remaining 6 percent combine both household and individual mobility.

Figure 3 shows the distribution of completed interviews per POU, decomposed into the contribution by stayers and movers. Each bar represents the share of total of POUs with *N* completed interviews; its height is a graphical representation of the bottom row of table 3. Within each category, the bottom segment of the bar represents the share of total POUs that came from stayers while the top 2 segments represent the contribution of movers.

Of POUs with 4 or fewer completed observations, movers account for 55 percent of the total. The share of movers drops substantially for those with 5 to 8 completed interviews, accounting for 27 percent on average; movers' share declines monotonically to zero. In-movers account for almost three-quarters of POUs with 1-3 completed interviews. There are about the same number of in-movers and out-movers with 4 completed interviews.

The significant decrease in the share of movers with more than 4 completed interviews is sensible. Even if the probability of moving stays constant, the greatest likelihood of observing a move lies in the 8 months when the person is not in the sample. This predicts a substantial, discrete fall in the share of movers after the first group of four months.

3.2 Demographic Characteristics of Movers and Stayers

Before assessing the bias from geographic mobility, this section examines characteristics of the populations of stayers and movers. If the population of persons that move is similar to those that do not, then their movement into and out of the CPS sample will cause little bias. However if the population of movers differs substantially from those who do not move, the bias from mobility may be large. In addition, it is important to distinguish between in-movers and out-movers. Even if movers differ from stayers, if persons who move into the CPS sample resemble those who leave it, then the bias from mobility may be small.

Table 4 reports the population proportions for several demographic characteristics.[20] The first column shows the proportion of all persons in the LPD with the indicated characteristic. The second column reports the proportion for stayers and the third and forth columns report the proportions for out-movers and in-movers.

The population of movers does not differ significantly in sex from those that do not move. Also, the populations of in-movers and out-movers have nearly the same ratio of females to males as the population of stayers. The other demographic characteristics have more meaningful differences. There are more

20. The populations are calculated using validity-weighted sampling weights; see section 2.4 for details.

nonwhite movers than stayers: the population of stayers is 85.1 percent white, compared with 81.8 percent for movers. Roughly 60 percent of the difference is accounted for by black movers. In-movers and out-movers do not differ appreciably in race.

A well-known feature of geographic mobility is the so-called "age selectivity of migration," which identifies a decline in mobility with age.[21] To assess this difference I classify age into 3 functional groups: younger (sixteen to twenty-four), prime age (twenty-five to fifty-four), and older (fifty-five and older). Table 4 confirms the age selectivity of migration: movers are younger than stayers. The population of movers has twice as many persons aged sixteen to twenty-four compared to stayers. Again, the difference between in-movers and out-movers is not large. The proportion of prime-age movers is basically the same as for stayers, implying an equally dramatic difference in the share of those aged fifty-five and older. The proportion of older movers is less than one-half that of stayers. Because prime-age workers are more likely to be in the labor force relative to those younger or older, the relative homogeneity in this category may mitigate potential bias from geographic mobility.

There are almost 80 percent more persons who have never married in the population of movers compared with stayers. Those never married account for 37 percent of movers but only 22 percent of stayers. The share of widowed and divorced are nearly identical between movers and stayers, implying that married persons are significantly less likely to move. The proportion married is 63 percent among stayers compared to 46 percent for movers.

There is relatively little difference in education between movers and stayers. The bottom panel of table 4 reports the distribution of educational attainment, divided into 4 functional categories: less than a high school education, high school graduates, some college, and college graduates. Movers are slightly more likely to be high school drop-outs or in college.

Although there is little or no difference in the distribution of movers and stayers by sex, race, and education, there are large differences in age and marital status. Individuals who move are more likely to be nonwhite, young, not married, and in college. In addition, because movers represent roughly 25 percent of all POUs, these differences will be economically meaningful if the characteristics are correlated with labor force status.

21. Gallaway (1969); Schlottmann and Herzog Jr. (1984); Tucker and Urton (1987); Peracchi and Welch (1994).

3.3 Labor Force Characteristics of Movers and Stayers

There are clear differences in the demographic characteristics of individuals who move and those who do not. This section explores whether those differences are also reflected in labor force status and transitions. As before, the first column of table 5 reports the proportion of all persons in the LPD with the indicated characteristic, the second column reports the proportion for stayers, and the third and forth columns report the proportions for out-movers and in-movers.

There are substantial differences in the distribution of labor force status between the movers and stayers (top panel, table 5). The population of movers has about one-fifth as many persons not in the labor force (NILF) and correspondingly more employed and unemployed. In particular, the there are twice as many more unemployed movers than stayers. Unlike with demographics, there are significant differences in labor force status between in-movers and out-movers.[22] There are about 7 percent more unemployed out-movers than in-movers, suggesting a link between job loss and mobility.

The lower panel of table 5 reports the population proportions for labor force transitions. Nontransitions, that is a "transition" between the same labor force state, are not reported.[23] The bottom 3 rows show unobserved transitions: transitions for which the previous month's labor force status is not known. These represent a substantial fraction of all transitions (30 percent). The discrepancy between measured stocks and gross flows that arises because of these unobservable transitions is known as "margin error."[24]

The first row shows that separations to unemployment (EU transitions) account for 0.62 percent of all labor force transitions in the CPS over 1976–2007. Among movers, however, EU transitions account for 0.93 percent of transitions, over 70 percent more than among stayers. Similarly, UE transitions account for an 80-percent larger share of mover's transitions than stayers'. Transitions between employment and nonparticipation also occur with greater frequency among movers, but the differences are more modest.

Most missing observations arise because of the CPS's rotating sample design, which ensures that at most 75 percent of the sample matches from one month to the next. However, unmatched observations also occur because of type A noninterviews, clerical errors, and mobility. Movers will have a greater

22. For similar findings see Bartel (1979); Schlottmann and Herzog Jr. (1981, 1984).

23. Nontransitions account for 66 percent of all transitions and 93 percent of *observed* transitions.

24. See Abowd and Zellner (1985); Poterba and Summers (1984, 1986); Chua and Fuller (1987); Fujita and Ramey (2006).

share of missing observations because out-movers are not followed and because the history of in-movers is unknown. In particular, because transitions are defined with respect to the current month, there are more unobservable transitions for in-movers than for out-movers.

This is confirmed in the bottom 3 rows of table 5, where movers have higher population proportions than stayers. In-movers record roughly 20 percent more missing transitions than do out-movers. A truly striking result is that transitions from missing to unemployment (`XU`) are almost three times as prevalent for movers. In contrast, transitions to employment are "only" 40 percent higher among movers. An important implication of these findings is that margin error–adjustment should be calculated separately for movers and stayers.[25]

3.4 Bias from Geographic Mobility

The population of individuals who move is different from those who do not.[26] Although the LPD identifies individuals that move and contains information on those persons while they are in the sample, it does not, of course, say anything about them when they are not in the sample.

Because the CPS does not follow households that move, estimates of movers' gross flows and hazard rates from the LPD may not accurately reflect that population's true behavior. However it is possible to conduct the counterfactual experiment of what the CPS data would show if there was no mobility by considering only the population of stayers. Comparing this counterfactual series with the actual series estimated from the entire population provides a bound on the bias from geographic mobility.

Let σ_t be the share of the month t population that does not move:

$$\sigma_t = \frac{\mathtt{P}_t^S}{\mathtt{P}_t}, \tag{2}$$

where $\mathtt{P}_t$ is the total population and superscript S denotes stayers. The average of σ_t over 1976–2007 is 0.7798. This value is not strictly comparable to the estimates of mobility rates presented earlier, which measure the number of persons not living at the same address one year later.

The total number of persons who transition from state $\mathtt{I}$ in month $t-1$ to state $\mathtt{J}$ in month t can be divided into the number of transitions made by stayers and that by movers:

$$\mathtt{IJ}_t = \mathtt{IJ}_t^S + \mathtt{IJ}_t^M, \tag{3}$$

25. See appendix B.

26. Whether these observed differences are the ex ante cause of mobility or the ex post result of mobility is a separate and interesting question.

where superscript M denotes movers. This implies a similar decomposition of the separation and job finding hazard rates:

$$s_t = \sigma_t s_t^S + (1-\sigma_t)s_t^M \quad \text{and} \quad f_t = \sigma_t f_t^S + (1-\sigma_t)f_t^M, \tag{4}$$

where s_t and f_t are the separation and job finding hazard rates for the entire CPS sample.

The monthly separation and job finding hazard rates are calculated by

$$\widehat{s}_t = \frac{\mathtt{EU}_t}{\mathtt{E}_{t-1}} \quad \text{and} \quad \widehat{f}_t = \frac{\mathtt{UE}_t}{\mathtt{U}_{t-1}} \tag{5}$$

for the entire CPS population and by

$$\widehat{s}_t^S = \frac{\mathtt{EU}_t^S}{\mathtt{E}_{t-1}^S} \quad \text{and} \quad \widehat{f}_t^S = \frac{\mathtt{UE}_t^S}{\mathtt{U}_{t-1}^S} \tag{6}$$

for the population of stayers, where $\mathtt{E}$ and $\mathtt{U}$ are the stock of employed and unemployed persons.

A way to assess the potential bias from geographic mobility is to measure the difference between the hazard rate calculated for stayers and the entire CPS sample. Define the ratio between the counterfactual hazard rate and the measured hazard rate as

$$G(s)_t = \frac{\widehat{s}_t^S}{\widehat{s}_t} \qquad \text{and} \qquad G(f)_t = \frac{\widehat{f}_t^S}{\widehat{f}_t} \tag{7}$$

If the hazard rates of the populations of movers and stayers are identical this ratio is 1; $G \neq 1$ indicates differences attributable to geographic mobility.

The upper panel of table 6 reports the averages of $G(s)$ and $G(f)$ over 1976–2007. The average ratio of the job finding hazard rate of movers to that of stayers is nearly 1, indicating that the job finding hazard rate of stayers does not differ much from that of the whole population. The average job finding hazard rate is about 2 percent lower for stayers than for the entire population.

In contrast, the separation hazard rate of stayers is almost 20 percent lower than that for the entire population. This implies that the separation rate for movers is much higher than in the total population. Indeed, the separation rate calculated from the available information from the population of movers is 65 percent higher than that from the entire sample. This value should be interpreted cautiously, however, because some labor force behavior of movers is not observable.

Nevertheless, there is a clear difference between movers and stayers in their probability of separating to unemployment: movers have a substantially higher separation hazard rate. Movers and stayers to do not differ significantly in job finding behavior, however. Although the effect on the level of separations is large, of principal concern is whether geographic mobility affects the *cyclical* behavior of hazard rates. If the difference between the separation rate of stayers and the general population does not change significantly over the business cycle, geographic mobility contributes little bias.

3.5 Cyclical Bias

I model the observed time series as the sum of four independent, unobserved components: a trend, a cycle, a seasonal, and an irregular component.[27] The trend represents low-frequency movements in the series. The cyclical component is a stochastic periodic function of time with a frequency at that of the business cycle. The seasonal component represents fluctuations that repeat annually and the irregular component captures the remaining non-systematic variation.

The structural time series model for the natural logarithm of each series, denoted y_t, is

$$y_t = \mu_t + \psi_t + \gamma_t + \varepsilon_t\,, \tag{8}$$

where μ_t is the trend, ψ_t the cyclical, γ_t the seasonal, and ε_t the irregular component. Details of the econometric specification of the components are provided in appendix B.

Equation 8 is recast as a state space model where the unobserved components are represented by the state of the system. The unknown parameters are estimated by maximum likelihood using the Kalman filter to update and smooth the unobserved state. The estimation is performed using Koopman et al. (2007)'s structural time series analyzer, modeller, and predictor (STAMP) program. See appendix B for details.

A reasonable concern is that mobility associated with the business cycle may create a cyclical bias in measured gross flows and hazard rates. First, however, it is important to understand how mobility changes over the business cycle. I measure the annual rate of geographic mobility by one minus the share of persons reported living at the same address one year later reported by the U.S. Census Bureau.[28]

27. This follows the general method described in Harvey (1989).
28. See U.S. Census Bureau (2007).

I isolate the cyclical component of mobility by estimating equation 8 at annual frequency.[29] Figure 4 plots the cyclical component of the mobility rate together with that of the unemployment rate for comparison.[30] The cyclical component of mobility tends to follow the unemployment rate, indicating that more people move during recessions than during booms. This is consistent geographic mobility as a means for reallocating idle labor to more productive uses. The contemporaneous correlation of the cyclical component of the mobility rate with the unemployment rate is 0.50, confirming the apparently countercyclicality. The peak correlation of 0.51 trails unemployment by two months.

I next estimate equation 8 for each of the four hazard rates in equations 5 and 6. I evaluate the cyclical bias using the ratio measure G (equation 7). In this case the ratio is calculated as the difference in the log cyclical components:

$$G(\psi^s)_t = \widehat{\psi}_t^{s,S} - \widehat{\psi}_t^s \qquad \text{and} \qquad G(\psi^f)_t = \widehat{\psi}_t^{f,S} - \widehat{\psi}_t^f, \tag{9}$$

where ψ^s and ψ^f are the cyclical components of the separation and job finding hazard rates calculated from the whole population and $\psi^{s,S}$ and $\psi^{f,S}$ are those calculated from only the population of stayers.

Summary statistics for $G(\psi^s)$ and $G(\psi^f)$ over 1976–2007 are reported in the lower panel of table 6. The values in the lower panel are the percentage difference between the cyclical component of the hazard rate calculated from the population of stayers and the hazard rate calculated from the entire population. The minimum and maximum values indicate that the greatest degree of bias from geographic mobility. The cyclical dynamics of job finding are not effected by mobility; the largest cyclical difference is 1 percent. There is a more modest effect of mobility on the separation hazard rate, although the peak bias never exceeds 4 percent.

Figures 5 and 6 plot the cyclical components of the actual and counterfactual hazard rate series (ψ_t and ψ_t^S). The cyclical component of the unemployment rate (in gray) is also shown for comparison. The solid line plots the hazard rate for the entire population while the dashed line uses only the population of stayers. The lower panel shows the estimated cyclical bias, $G(\psi)_t$. The vertical axes are drawn so the divisions of the left and right ordinates have the same size.

The separation hazard rate of stayers, shown in figure 5, is more volatile at business cycle frequencies than the separation hazard rate of the entire population. It generally falls further at the cyclical peak (the trough of unemployment)

29. This eliminates the seasonal component.

30. For the graph only, I use a locally weighted polynomial regression smoother (Cleveland, 1979) to create a monthly time series of the cyclical component from the annual data.

and rises higher at the cyclical trough (the peak of unemployment). The cyclical bias, shown in the lower panel of figure 5, reflects this pattern. The cyclical correlation of the bias with the unemployment rate (table 7) is 0.55, indicating moderate countercyclicality. That is, the bias from geographic mobility rises during recession as more people move.

Figure 6 shows that there is little effect of geographic mobility on job finding hazard rates. The hazard rate calculated from the population of stayers is largely indistinguishable from that calculated using the entire population. Although the bias from geographic mobility, shown in the lower panel of figure 6, is mildly procyclical (table 7), the difference between job finding hazard rates measured from stayers and the whole population never exceeds 1 percent.

3.6 Discussion

The LPD allows me to identify individuals who move into and out of the CPS sample. Because many movers spend four months or more in the sample, I can observe their demographic characteristics and establish a meaningful history of labor market behavior. Comparing the populations of movers and stayers reveals no difference in the composition of sex and minor differences in race and education. There are, however, large differences in age and marital status of movers compared with stayers.

There are also substantial differences in the distribution of labor force status between the two populations: there are 60 percent more unemployed movers than unemployed stayers. In addition, EU separations and UE accessions comprise almost twice the share of transitions for movers than for stayers. Separations and accessions are best interpreted in the context of separation and job finding hazard rates. Movers have a substantially higher separation hazard rate than stayers, although they to do not differ significantly in job finding rate.

Geographic mobility varies negatively with the business cycle, possibly creating a cyclical bias to measured separation and job finding hazard rates. The bias in hazard rates arising from not observing the behavior of movers can be assessed by comparing a counterfactual hazard rate calculated from the population of stayers to the hazard rate calculated for the entire population.

The cyclical bias in the separation hazard rate is countercyclical, meaning that the separation hazard rate calculated using the entire CPS sample will appear too *a*cyclical. There is little effect of geographic mobility on the job finding hazard rate.

This evidence can be interpreted as follows. The rate of separations to unemployment and of geographic mobility both increase during a recession. The separation hazard rate of stayers rises more during a recession than does the entire sample, implying that the separation rate of movers is *less* countercyclical.[31] Put differently, during a boom the separation hazard rate falls, however the separation rate of movers falls by less than the entire population. Nevertheless, the cyclical difference between the separation hazard rate of stayers and the entire population never exceeds 4 percent; geographic mobility does not significantly affect the cyclicality of measured hazard rates.

This relatively small bias seems at odds with the substantial differences in average separation hazard rates between movers and stayers (table 6). These differences can be reconciled by recognizing the importance of differentiating between out-movers and in-movers. The argument of geographic mobility bias is one of sample attrition: a person leaves the sample and is not followed. But focusing solely on out-movers is misguided. There are equally as many in-movers as out-movers (by person) and in-movers account for 60 of movers' observations.

In addition, the demographic and labor force evidence presented in tables 4 and 5 shows that, although movers are quite different from stayers, the differences between in-movers and out-movers are small, especially relative to stayers. Thus, appealing to the CPS's random sampling, a person who moves out of one address is replaced by a similar in-mover elsewhere in the country and the true bias from sample attrition (i. e., out-movers) is offset by similar in-movers.

4 Robustness

This section examines the robustness of the analysis of geographic mobility. I consider alternative measures of validating matches of observations in the LPD and assess my findings using an alternate procedure for isolating cyclical components.

4.1 Alternate Measures of Match Validity

This section evaluates match validity using 2 alternative schemes. The first scheme is "naive" matching, that is matches determined solely by the information that defines which observations can match. A second scheme is to consider

31. This is confirmed by estimating the cyclical component of movers' separation hazard rate. The cyclical correlation of with unemployment is 0.85, compared with 0.88 in the entire sample. As before, this relationship should be interpreted cautiously because the full history of the population of movers is not observed.

the average validity score for a person,

$$\overline{V}_i = \frac{1}{T_i} \sum_{t=1}^{T_i} V_{it}, \tag{10}$$

where V_{it} is defined in equation 1, and use a threshold rule to determine which matches are counted. I included all persons where $\overline{V}_i \geq 0.875$. In practice, this value allows for 1 failure among the 3 criteria over a four month block of observations.

I then calculate the separation and job finding hazard rates under each of the alternate schemes. To facilitate comparison across the 3 schemes, the hazard rates are expressed relative to the baseline scheme (weighted matching). Table 8 reports the summary statistics for these two measures. There is virtually no difference in the measures. Naive matching yields almost identical results as probability-weighted matching. There are larger differences when using the threshold criterion, but the effects are still quite modest. Both separation and job finding hazard rates are slightly lower, with a peak difference of about 4 percent.

There are two central results. First, overall match quality in the LPD is very high. The average V over all observations in the LPD is 0.9930. This is not a feature of the LPD per se, but of the underlying CPS data. The second is not surprising given the first: adjusting for matches supplemental validity does not significantly affect results.

4.2 Alternate Method of Isolating Cyclical Component

In this section I explore an alternate method for isolating the cyclical component of the time series. A common technique in macroeconomics is to filter the seasonally-adjusted series using the Hodrick-Prescott (HP) filter to extract the cyclical component.[32]

Because the HP filter requires a continuous time series, any missing observations associated with changes in the household identifier must be interpolated. Researchers use either a local moving average or linear interpolation to create a continuous time series. For simplicity, I use linear interpolation. I next seasonally adjust the series using the Census Bureau's X-12-ARIMA seasonal adjustment program. Finally, I HP filter the seasonally-adjusted series with smoothing parameter $\lambda = 129,600$.[33]

32. Hodrick and Prescott (1997).

33. Ravn and Uhlig (2002) find the optimal HP smoothing parameter for monthly data is 129, 600.

Figures 7 and 8 plot the cyclical components of the separation and job finding hazard rate using the alternative cyclical isolation procedure. The cyclical components are considerably more volatile, particularly at high frequencies. This high-frequency volatility is a natural consequence of the HP filter, which removes only the low-frequency trend. As such, it is more difficult to clearly identify cyclical patterns from the graph of the time series than in figures 5 and 6, particularly distinguishing between the actual and counterfactual series.

As the lower panels in figures 7 and 8 show, the cyclical bias estimated using the HP filter is considerably larger than that estimated using the unobserved-components model. The bias from geographic mobility contributes up to 15 percent for the separation rate and 10 percent for the job finding rate (table 9).

Although the degree bias is larger, its correlation with unemployment weakens dramatically when using the HP filter (table 10). Although the signs remain the same, the correlations in the HP data are essentially acyclical. The correlation of the bias in the job finding hazard rate with unemployment is not statistically significant at the 10 percent level. Even though the degree of bias is large, it is unrelated to the business cycle identified by the HP filter.

Note also that the cyclical correlation of the hazard rates fall when using the HP filter to isolate the cyclical component. Using data for the entire population, the cyclical correlation with unemployment for the separation hazard rate falls from 0.87 to 0.60 and from -0.94 to -0.73 for the job finding hazard rate. Nevertheless the HP filter shows a strong relationship of both hazard rates with the business cycle.

5 Conclusion

Because the CPS does not follow individuals that move away from a sample address, the strong relationship between unemployment, job separation, and mobility creates concern that labor market dynamics captured by the CPS may be biased from sample attrition related to geographic mobility. Using a new database that permits sophisticated longitudinal analysis of the all CPS data, I find that the cyclical bias arising from geographic mobility is small. At business cycle frequencies, the difference between the separation hazard rate calculated from the entire CPS sample and from a subset that are known not to have moved never exceeds 4 percent. There is little effect from geographic mobility on the job finding hazard rate.

To facilitate this study of mobility and of other important longitudinal research topics, I construct a new database, the Longitudinal Population Database

(LPD), that organizes the CPS data into individual panels, where the person is the fundamental unit. I develop a novel framework for identifying an individual's full longitudinal history inside a survey that is, fundamentally, a sample of addresses. The LPD, contains the complete interview history for every person surveyed by the CPS over 1976–2006, over 10 million individuals. Over 65 percent of persons have a interview history of at least four continuous months and about 4.5 million have a complete history of 8 observations.

The LPD provides excellent information on mobility. Because the LPD contains the entire history of each address in the sample, it is possible to distinguish between movers and stayers and between in-movers and out-movers. About 25 percent of individuals in the LPD move at some point during their tenure in the sample. Since many movers spend at least four months in the sample, the LPD records their demographic characteristics and a meaningful history of labor force behavior. Furthermore, because the selection of an address for sampling is independent of the decision to move, the LPD contains a true random sample of movers.

Comparing the populations of movers and stayers reveals only minor difference between in the composition of sex, race, and education of movers and stayers. However, I find that movers are younger than stayers and more movers are unmarried. Movers are also more likely to be unemployed.

I assess labor market dynamics using the separation and job finding hazard rates. On average, the separation hazard rate of stayers is almost 20 percent lower than that for the entire population, implying a high separation rate for movers. The separation rate of movers is, indeed, about 65 percent higher than when using the entire population. There is relatively little difference in the job finding hazard rates between movers and stayers.

This large difference in average separation hazard rates between movers and stayers seems at odds with the small degree of cyclical bias. This tension can be reconciled by distinguishing between out-movers and in-movers. The argument of geographic mobility bias is one of sample attrition: a person leaves the sample and is not followed. But focusing solely on out-movers is misguided because there are equally as many in-movers as out-movers. In addition, the demographic and labor force evidence shows that the differences between in-movers and out-movers are small relative to those between movers and stayers.

The evidence presented in this paper is consistent with the idea that geographic mobility reflects efficient resource reallocation. Geographic mobility increases during a recession, facilitating the reallocation of idle resources—unemployed persons—across space to more productive uses. Fortunately, this labor reallocation does not significantly impact the measurement of U.S. labor market dynamics.

A Appendix A

The appendix provides details about the construction of the LPD. The database is compiled in two stages. In the first stage the raw data for each month are imported into a statistical program and processed to ensure that all variables are longitudinally consistent across all marks. In the second stage the processed monthly files are appended together to create a longitudinal data set. The entire data set is then processed to properly identify addresses, households and household changes, and individuals using all longitudinal information.

A.1 Stage I: Raw Data

In stage I, the monthly data files are processed individually. Each is imported into a statistical program and then processed to create longitudinally-consistent variables.

The monthly public-use CPS microdata flat files are downloaded from the Census Bureau and the CPS data repository at the National Bureau of Economic Research (NBER). The Census Bureau web site hosts the microdata files for 1992 to the present. Data before 1992 come from the NBER, which maintains copies of the data files for 1976 to the present.

The variable layout and definitions in the microdata files change 17 times over 1976–2006. Each different version of the layout and definition is called a "mark." Many of the variable locations and definitions remain the same across marks, however a change in any one variable constitutes a new mark. The table below lists the 18 marks and the months they span.

Mark number	*Start date*	*End date*	*No. of months*
0	Jan 1976	Dec 1977	24
1	Jan 1978	Dec 1981	48
2	Jan 1982	Dec 1982	12
3	Jan 1983	Dec 1983	12
4	Jan 1984	Jun 1985	18
5	Jul 1985	Dec 1985	6
6	Jan 1986	Dec 1988	36
7	Jan 1989	Dec 1991	36
8	Jan 1992	Dec 1993	24
9	Jan 1994	Mar 1994	3
10	Apr 1994	May 1995	14
11	Jun 1995	Aug 1995	3

Mark number	*Start date*	*End date*	*No. of months*
12	Sep 1995	Dec 1997	28
13	Jan 1998	Dec 2002	60
14	Jan 2003	Apr 2004	16
15	May 2004	Jul 2005	15
16	Aug 2005	Dec 2006	17
17	Jan 2007	Dec 2007	12

A.2 Stage I: Data Dictionaries

To construct a longitudinal database, all variable names and definitions must be the same across all marks. Because many change from mark to mark, I compare the 18 data definition files and create a set of universal variable names and definitions that are consistent across all marks. Table 11 reports the universal variable names and definitions.

I then create dictionaries for each mark that correspond to the universal definitions. Jean Roth at the NBER provides data dictionaries for marks 7–16, however they do not conform to the universal definitions. I modify these dictionaries to maintain longitudinal consistency.

A.3 Stage I: Longitudinal Consistency

This subsection describes how the LPD's variables are created from CPS variables to ensure longitudinal consistency across all marks. Variable names are set in monospaced type; those in uppercase identify variables from the CPS while those in lowercase are LPD variables.

A.3.1 Survey Date

All observations in the CPS contain the 2-digit month of the survey (`HRMONTH`) and some measure of the year (`HRYEAR`). For marks 1–8 the CPS reports only the last digit of the survey year, while marks 9–12 report the last 2 digits. All other marks include the 4-digit year. The LPD variable year is constructed from `HRYEAR` to report the full 4-digit year of survey. The information on month is unaltered.

A.3.2 Interview Status

The CPS reports the status of each interview in the variable `HRINTSTA`. The interview status can take on 4 values: completed interview, type A noninterview, type B noninterview, and type C noninterview. The LPD variable `INTSTAT` reports this code for marks 9–16. Prior to mark 9 the CPS classifies interview status into only 3 categories, combining type B and type C noninterviews into one category. For these marks the type B and type C noninterviews are separated using supplementary information.

A.3.3 State

The CPS records the U.S. state of the address using two different code systems. Marks 7–16 report the state using both the Federal Information Processing System (FIPS) code (`GESTFIPS`) and the Census Bureau state code (`GESTCEN`). Prior to mark 6, the CPS reports only `GESTCEN`. The LPD variable `STATE` contains the FIPS state code for the address. The concordance between Census state codes and FIPS state codes is below.

State	*FIPS code*	*Census code*	*State*	*FIPS code*	*Census code*	*State*	*FIPS code*	*Census code*
AK	02	94	KY	21	61	NY	36	21
AL	01	63	LA	22	72	OH	39	31
AR	05	71	MA	25	14	OK	40	73
AZ	04	86	MD	24	52	OR	41	92
CA	06	93	ME	23	11	PA	42	23
CO	08	84	MI	26	34	RI	44	15
CT	09	16	MN	27	41	SC	45	57
DC	11	53	MO	29	43	SD	46	45
DE	10	51	MS	28	64	TN	47	62
FL	12	59	MT	30	81	TX	48	74
GA	13	58	NC	37	56	UT	49	87
HI	15	95	ND	38	44	VA	51	54
IA	19	42	NE	31	46	VT	50	13
ID	16	82	NH	33	12	WA	53	91
IL	17	33	NJ	34	22	WI	55	35
IN	18	32	NM	35	85	WV	54	55
KS	20	47	NV	32	88	WY	56	83

A.3.4 Sex

No changes to the coding are required.

A.3.5 Race

The level of detail for racial classification varies widely across the marks. There are 3 major classification schemes. The most recent marks (14–16) classify race into 21 separate categories (PRDTRACE). Marks 7–13 have 5 distinct categories (PERACE) and marks 1–6 report only 3: white, black, and other. Thus, to maintain longitudinal consistency, race is recoded into the 3 categories from marks 1–6. Below is a concordance for the two other schemes.

Mark 14–16		*Mark 7–13*	
PRDTRACE	RACE	PERACE	RACE
WHITE	WHITE	WHITE	WHITE
BLACK	BLACK	BLACK	BLACK
AMERICAN INDIAN (AI)	OTHER	AMERICAN INDIAN	OTHER
ASIAN	OTHER	ASIAN-PACIFIC ISLANDER	OTHER
HAWAIIAN (HP)	OTHER	OTHER	OTHER
WHITE-BLACK	WHITE		
WHITE-AI	WHITE		
WHITE-ASIAN	WHITE		
WHITE-HP	WHITE		
BLACK-AI	BLACK		
BLACK-ASIAN	BLACK		
BLACK-HP	BLACK		
AI-ASIAN	OTHER		
ASIAN-HP	OTHER		
WHITE-BLACK-AI	WHITE		
WHITE-BLACK-ASIAN	WHITE		
WHITE-AI-ASIAN	WHITE		
WHITE-ASIAN-HP	WHITE		
WHITE-BLACK-AI-ASIAN	WHITE		
2 OR 3 RACES	OTHER		
4 OR 5 RACES	OTHER		

A.3.6 Age

The CPS reports each individual's age as of the end of the reference week (PEAGE), topcoded at different years depending on the mark. For most variables the CPS reports information with greater detail as the survey ages, but this is not the case with age: marks 1–4 topcode ages above 99 years old, marks 5–14 topcode ages above 90, and marks 15–16 topcode ages above 80. The LPD variable AGE is re-topcoded as 80 for ages 80–84 and as 85 for ages 85 and older.

A.3.7 Marital Status

The CPS classifies marital status (PEMARITL) using 3 different schemes. The LPD classifies marital status (MS) as either married, widowed/divorced, or never married. The concordance with the CPS data is below.

PEMARITL	MS
Mark 9–16	
MARRIED-SPOUSE PRESENT	MARRIED
MARRIED-SPOUSE ABSENT	MARRIED
WIDOWED	WIDOWED/DIVORCED
DIVORCED	WIDOWED/DIVORCED
SEPARATED	MARRIED
NEVER MARRIED	NEVER MARRIED
Mark 6–8	
MARRIED-CIVILIAN SPOUSE PRESENT	MARRIED
MARRIED-AF SPOUSE PRESENT	MARRIED
MARRIED-SPOUSE ABSENT	MARRIED
WIDOWED	WIDOWED/DIVORCED
DIVORCED	WIDOWED/DIVORCED
SEPARATED	MARRIED
NEVER MARRIED	NEVER MARRIED
Mark 1–5	
MARRIED-CIVILIAN SPOUSE PRESENT	MARRIED
MARRIED-AF SPOUSE PRESENT	MARRIED
MARRIED-SPOUSE ABSENT	MARRIED
WIDOWED OR DIVORCED	WIDOWED/DIVORCED
NEVER MARRIED	NEVER MARRIED

A.3.8 Educational Attainment

As part of the 1994 survey redesign, the CPS changed the education question from a quantitative question about the years of schooling attended to a qualitative question about level of education attained. Jaeger (1997) studied the relationship between the two questions by comparing responses from individuals who answered both versions of the question. The LPD education variable, EDUC, is coded using Jaeger's correspondence, reported below.

	Highest grade attended		*Educational*
Category	*Not completed*	*Completed*	*attainment*
High school dropout	0–12	1–11	31–37
High school graduate	n.a.	12	38, 39
Some college	13–16	13–15	40–42
College graduates	17, 18	16–18	43–46

A.3.9 Labor Force Status

The labor force status ultimately reported in the CPS (PEMLR) is a recode based on answers to survey questions. Although the broad classification of labor force status—employed, unemployed, and NILF—is unchanged throughout the history of the CPS, the labor force subclassifications do change. The LPD classifies labor force status into those 3 broad categories as follows:

PEMLR	LFS
Mark 12–16	
EMPLOYED-AT WORK	EMPLOYED
EMPLOYED-ABSENT	EMPLOYED
UNEMPLOYED-ON LAYOFF	UNEMPLOYED
UNEMPLOYED-LOOKING	UNEMPLOYED
NILF-RETIRED	NILF
NILF-DISABLED	NILF
NILF-OTHER	NILF
Mark 6–11	
EMPLOYED-AT WORK	EMPLOYED
EMPLOYED-ABSENT	EMPLOYED
UNEMPLOYED-ON LAYOFF	UNEMPLOYED
UNEMPLOYED-LOOKING	UNEMPLOYED
NILF-WORK W/O PAY	NILF
NILF-UNAVAILABLE	NILF
NILF-OTHER	NILF
Mark 1–5	
EMPLOYED-AT WORK	EMPLOYED
EMPLOYED-ABSENT	EMPLOYED
UNEMPLOYED-LOOKING	UNEMPLOYED
NILF-HOUSE	NILF
NILF-SCHOOL	NILF
NILF-UNABLE	NILF
NILF-OTHER (INC. RETIRED)	NILF

A.3.10 Industry and Occupation

No changes to the coding are required.

A.4 Stage II: Observation Identifiers

The first step of stage II is to append the processed monthly data files together into a single longitudinal data set. The observations are then sorted chronologically by a unique address identifier to organize them into a time series for each address. The addresses are then processed to identify households and the households processed to identify individuals. The following sections describe how these identifiers are constructed.

A.4.1 Address Identifier

The primary identifier in the CPS is a "unique household identifier" (`HRHHID`). All observations in the CPS have a `HRHHID`. This unfortunately-named variable does not, in fact, identify households nor is it unique, either locally (within a single month) or globally (both within and across months). More precisely, it is a partial *address* identifier that, together with other variables, uniquely identifies an address. For marks 1–11 `HRHHID` is a 12-digit number, but for marks 12–16 it increases to 15 digits. All 12-digit `HRHHID`s are padded to 15 digits by adding 3 leading zeros.

For mark 10 and later, an address is uniquely identified by `HRHHID` and 2 other variables: the sample identifier (`HRSAMPLE`) and serial suffix (`HRSERSUF`). Concatenating these three variables creates a 19-digit, globally-unique address identifier, `AID`. An AOU is defined technically as all observations with the same `AID`.

Unfortunately, the 2 additional variables needed to create `AID` do not exist in marks 1–10. This problem manifests itself when the data from all months are combined into one longitudinal data set. Because an address is not uniquely identified across months, observations from several different addresses will have the same `HRHHID`. This collection of all observations from a single household identifier is called a `HRHHID` group. The figure below illustrates the problem. The top row displays fictional data for an address uniquely identified by `HRHHID` (and the supplemental variables), such as one from marks 11–16. The cell values are the address's month in sample (MIS). The bottom two rows show fictional data for 2 `HRHHID` groups from marks 1–10. Each `HRHHID` has observations for many more months than is possible under the CPS survey design.

	Month																									
HRHHID	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**	**24**	**24**	**26**
…0026						1	2	3	4									5	6	7	8					
…5923	1	2	3	4	1	2	3	4	1	2	3	4	5	6	7	8	5	6	7	8	5	6	7	8	1	2
…8321	7	8	1	2	3	4	1	2	3	4	1	2	3	4	5	6	7	8	5	6	7	8	5	6	7	8

Because each address is surveyed according to a defined rotation pattern, there is a unique relationship between the survey date and an address's MIS within an `HRHHID` group. For example, an address that enters the CPS sample at calendar month 6 can have at most 4 interviews at months 6–9 and 4 more interviews at months 18–21 (for example, the `HRHHID` ending with 0026 in the figure above). If the data from the `HRHHID` group at month 18 does not have MIS = 5, then it must be from a different address. I have written an algorithm that exploits this relationship to uniquely identify individual addresses within a `HRHHID` group.

The figure below illustrates how the observations from the 2 fictional `HRHHID` groups are separated into different addresses under the address algorithm.

	Month																									
HRHHID	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**	**24**	**24**	**26**
…5923	1	2	3	4	1	2	3	4	1	2	3	4	5	6	7	8	5	6	7	8	5	6	7	8	1	2
…5923A	1	2	3	4									5	6	7	8										
…5923B					1	2	3	4									5	6	7	8						
…5923C									1	2	3	4									5	6	7	8		
…5923D																									1	2
…8321	7	8	1	2	3	4	1	2	3	4	1	2	3	4	5	6	7	8	5	6	7	8	5	6	7	8
…321A	7	8																								
…8321B			1	2	3	4									5	6	7	8								
…8321C							1	2	3	4									5	6	7	8				
…8321D											1	2	3	4									5	6	7	8

About 17 percent of AOUs have no completed interviews. Of these over 85 percent consist exclusively of type B or type C noninterviews. These AOUs are discarded because they contribute no data, longitudinal or otherwise.[34] This is different from discarding a single interview that has no data in a particular month; these AOUs are never eligible for interview during their entire CPS history. Line 4 in table 1 depicts an example of an AOU that would be discarded. The remaining AOUs with no completed interviews consist of all type A noninterviews. These AOUs remain in the sample because these addresses contain households that could have been interviewed.

34. If addresses are ineligible at random, excluding them does not bias the sample. If, however, a disproportionate number of addresses selected were located, for example, in a poor innercity and had been condemned, then excluding these AOUs could bias the estimate. Comparing the distribution across states of AOUs with noninterviews against those without noninterviews reveals no substantive differences.

A.4.2 Household Identifier

After creating a unique address identifier, all addresses are processed to identify unique households. Section 3.2 describes the 4 ways a household change can occur within an AOU. I have written an algorithm that identify these household changes and create a unique household identifier.

The CPS records the number of households that occupy an address during its 8-interview history. Each time a new household is identified at an address, the household number (HUHHNUM) is incremented. There may be up to 8 different households at an address. For addresses without a noninterview, the household number (HNUM) is given by HUHHNUM. This correctly identifies type H1 household changes (no intervening vacancy).

Addresses with noninterviews require special processing to create the correct household number. The CPS does not change HUHHNUM following a type B or type C noninterview.[35] Therefore, all observations after a type B or type C noninterview are assigned to the same household when they must be from a different household. The household algorithm correctly identifies the remaining 3 types of household change.

For each address, the household algorithm examines all observations with the same household number in chronological order. This is the largest group of observations that could be from the same household. When it encounters a type B or type C noninterview it does the following:

1. if there are no valid (completed interview or type A noninterview) observations in the past, the current observation is dropped;

2. if there are valid observations in the past but none in the future, the current observation is dropped; or

3. if there are valid observations in the past and valid observations in the future, the current observation is dropped and all future observations from this address are assigned the next HNUM.

The algorithm continues until all observations from an address with the same household number have been processed. It then repeats for the next address. Appending HNUM to AID creates a 20-digit, globally-unique household identifier, HID. A HOU is defined technically as all observations with the same HID.

35. Type B and type C noninterviews indicate the address is ineligible for interview that month, implying that the future (previous) occupants are not the same as the previous (future) occupants. Had the same occupants simply been unavailable that month, the interview would have been recorded as a type A noninterview.

A.4.3 Person Identifier

The CPS identifies individuals within a household by their line number on the survey response sheet (`PULINENO`). An individual retains the same line number for each month in the survey. Appending the 2-digit line number to hid creates a 22-digit, globally-unique person identifier, `PID`. When `PULINENO` is less than 2 digits, a leading zero is added. A POU is defined technically as all observations with the same `PID`.

This procedure does not work, however, for persons at an address with a type A noninterview.[36] Because no survey is performed that month, no information on the number of persons at the address is collected. Thus, because no line number exists that month, the longitudinal continuity of all POUs at the address is interrupted.

The third algorithm I create processes households for noninterviews and generates line numbers for persons living at the address in months with noninterviews. For each household, the person algorithm searches in chronological order for a noninterview. When it finds a noninterview it does the following:

1. if there is a valid observation in the previous month, the current month's observation is duplicated for each person at the address during the previous month;

2. if there are no valid observations in the past but there is a valid observation in the future, the current month's observation is duplicated for each person at the address during the first future month with a valid observation; or

3. if there are no valid observations in the past or in the future, the current month's observation is given a line number of 1.

The newly-created observations for each person contain the same information as the original address-level observation. That is, they have only information on the address and interview status. They contain no demographic or labor force information. The person algorithm does not currently attempt to impute this missing information.

It is impossible to know, a priori, if the persons who occupy the address in a month with a noninterview are the same as those who occupy the address in the previous or subsequent months. The algorithm assigns line numbers based on the last known observation or, when no previous valid observation exists, on the closest future observation.

36. Note that at this point in processing all type B and type C noninterviews have been dropped. The term noninterview thus refers only to type A noninterviews.

B Appendix B

B.1 Correcting for Margin Error

Because of the CPS's rotating sample design, at most 75 percent of observations can be matched across succeeding months. The simplest approach is to assume that the unmatched observations are simply missing at random; calculations are performed on the population of matched observations. This assumption has been shown to be a poor one.[37] In particular, the missing-at-random (MAR) correction significantly undercounts the unemployed.

The conditional MAR model is a simple but powerful extension of the MAR model. Given the timing convention for flows, a person's month t labor force status is always observed, even if the previous month's status is unknown. The MAR model throws this information away. Similar to the corrections of Abowd and Zellner (1985) and Fujita and Ramey (2006), the conditional MAR correction makes use of partially-classified observations. In particular, it assumes that a person missing in month $t-1$ with state $\mathtt{J}$ in month t is drawn randomly from the population of persons with state $\mathtt{J}$ in month t. That is, a person is missing at random conditional on having state $\mathtt{J}$ in month t.

The BLS performs much of its second-stage analysis separately by demographic group.[38] In particular, the distinction between male and female and between white and nonwhite are most important. In addition, this paper shows that the distinction between mover and stayer is of critical importance. I adjust for margin error separately by these sex, race, and mobility.

Let $m = 0$ for persons that do not move and $m = 1$ for movers. Let $\mathtt{IJ}_{srmt}$ be the number of persons with sex $s \in \{M, F\}$ and race $r \in \{W, NW\}$ who had labor force status $\mathtt{i}$ in month $t-1$ and status $\mathtt{j}$ in month t. Let $\mathtt{MJ}_{srmt}$ be the number of persons with missing labor force status in month $t-1$ and status $\mathtt{j}$ in month t. The ratio

$$R^{\mathtt{J}}_{srmt} = \frac{\mathtt{EJ}_{srmt} + \mathtt{NJ}_{srmt} + \mathtt{UJ}_{srmt}}{\mathtt{EJ}_{srmt} + \mathtt{NJ}_{srmt} + \mathtt{UJ}_{srmt} + \mathtt{MJ}_{srmt}} \tag{11}$$

is the number of observed transitions into state $\mathtt{J}$ (flows into $\mathtt{J}$) relative to the total number persons who had status $\mathtt{J}$ in month t (stock of $\mathtt{J}$). The MAR correction normalizes the entire population to the sum of all observed transitions in each state:

$$R^{MAR}_{srt} = \frac{\mathtt{IJ}_{srt}}{\sum_{\mathtt{i}\in\{\mathtt{E,N,U}\}} \sum_{\mathtt{j}\in\{\mathtt{E,N,U}\}} \mathtt{ij}_{srt}}. \tag{12}$$

37. See Abowd and Zellner (1985); Poterba and Summers (1986).
38. Bureau of Labor Statistics (2002).

Define the margin error–adjusted $\mathtt{IJ}$ flow, denoted with a tilde, for mobility m, sex s, and race r in month t as

$$\widetilde{\mathtt{IJ}}_{srmt} = \frac{\mathtt{IJ}_{srmt}}{R^{\mathrm{J}}_{srmt}}. \tag{13}$$

The aggregate margin error–adjusted $\mathtt{IJ}$ flow in month t is the sum over sex and race categories:

$$\widetilde{\mathtt{IJ}}_{t} = \sum_m \sum_s \sum_r \widetilde{\mathtt{IJ}}_{srmt}. \tag{14}$$

B.2 Structural Time Series Model

The structural time series model for the natural logarithm of each series, denoted y_t, is

$$y_t = \mu_t + \psi_t + \gamma_t + \varepsilon_t, \tag*{[8]}$$

where μ_t is the trend, ψ_t the cyclical, γ_t the seasonal, and ε_t the irregular component.

I model each trend component as a smooth first-order local linear trend:

$$\mu_t = \mu_{t-1} + \beta_{t-1} + \eta_t \tag{15}$$

$$\Delta\beta_t = \zeta_t, \tag{16}$$

where $\Delta = (1 - L)$ and L is the lag operator. The disturbances η_t and ζ_t are independent and identically distributed (i. i. d.) normal random variables with mean zero and variances σ^2_η and σ^2_ζ.

The cyclical component is modeled as a second-order stochastic cycle with frequency λ, where[39]

$$\begin{bmatrix} \psi_t^{(j)} \\ \psi_t^{*(j)} \end{bmatrix} = \rho \begin{bmatrix} \cos\lambda & \sin\lambda \\ -\sin\lambda & \cos\lambda \end{bmatrix} \begin{bmatrix} \psi_{t-1}^{(j)} \\ \psi_{t-1}^{*(j)} \end{bmatrix} + \begin{bmatrix} \psi_t^{(j-1)} \\ \psi_t^{*(j-1)} \end{bmatrix} \tag{17}$$

for $j = 1, 2$ and $\psi_t^{(0)} = \kappa_t$ and $\psi_t^{*(0)} = \kappa_t^*$. The disturbances κ_t and κ_t^* are i. i. d. normal each with mean zero and variance σ^2_κ. Note that for $j = 1$ and $\rho = 1$ equation 17 reduces to a deterministic cycle

$$\psi_t = \psi_0 \cos\lambda t + \psi_0^* \sin\lambda t,$$

39. Harvey and Trimbur (2003) find that, in practice, a second-order cycle provides a good approximation of the gain function of the Baxter-King (BK) bandpass filter.

where ψ_0 and ψ_0^* are i. i. d. zero-mean random variables with variance σ^2_ψ.

The stochastic seasonal component is constructed so that the s seasonal effects sum to zero in expectation. This is modeled as

$$\gamma_t = -\sum_{j=1}^{s-1} \gamma_{t-j} + \omega_t, \tag{18}$$

where $\omega_t \sim N(0, \sigma^2_\omega)$. Finally, the irregular component ε_t is i. i. d. normal with zero mean and variance σ^2_ε. All disturbances are mutually uncorrelated.

The model given by equations 8 and 15–18 is represented by the state space system relating observed data y_t to the unobserved state vector $\boldsymbol{\alpha}_t$ through a measurement vector $\mathbf{z}$:

$$y_t = \mathbf{z}'\boldsymbol{\alpha}_t + \varepsilon_t \tag{19}$$

$$\boldsymbol{\alpha}_t = \mathbf{T}\boldsymbol{\alpha}_{t-1} + \boldsymbol{\eta}_t. \tag{20}$$

The unobserved state evolves according to a first-order Markov process with transition matrix $\mathbf{T}$. The state equation (20) is

$$\begin{bmatrix} \mu_t \\ \beta_t \\ \psi_t \\ \psi_t^* \\ \gamma_{t-1} \\ \gamma_{t-2} \\ \vdots \\ \gamma_{t-s+2} \end{bmatrix} = \begin{bmatrix} \mathbf{T}_{\text{trend}} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{T}_{\text{cycle}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{T}_{\text{seasonal}} \end{bmatrix} \begin{bmatrix} \mu_{t-1} \\ \beta_{t-1} \\ \psi_{t-1} \\ \psi_{t-1}^* \\ \gamma_{t-2} \\ \gamma_{t-3} \\ \vdots \\ \gamma_{t-s+1} \end{bmatrix} + \begin{bmatrix} \eta_t \\ \zeta_t \\ \kappa_t \\ \kappa_t^* \\ \omega_t \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{21}$$

where

$$\mathbf{T}_{\text{trend}} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$

$$\mathbf{T}_{\text{cycle}} = \begin{bmatrix} \rho\cos\lambda & \rho\sin\lambda \\ -\rho\sin\lambda & \rho\cos\lambda \end{bmatrix}$$

$$\underset{(s-1\times s-1)}{\mathbf{T}_{\text{seasonal}}} = \begin{bmatrix} -1 & -1 & \dots & -1 & -1 \\ 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ & & \vdots & & \\ 0 & 0 & \dots & 1 & 0 \end{bmatrix}$$

This system represents a system with a first-order cycle. The extension to second-order cycles is straightforward.

The state vector enters the measurement equation by the $(4+s-1 \times 1)$ vector

$$z = \begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 0 & \dots & 0 \end{bmatrix}'. \tag{22}$$

The unknown parameters σ_ε^2, σ_η^2, ρ, λ, σ_κ^2, and σ_ω^2 are estimated by maximum likelihood using the Kalman filter. For consistency across all series, I fix the variance of the trend so as to reproduce the HP trend.[40] This variance is $\sigma_\zeta^2 = \sigma_\varepsilon^2/129,600$.[41] The cycle frequency λ is fixed at sixty months; this corresponds roughly with the center of Burns and Mitchell (1946)'s period of business cycle frequencies. With these restrictions, the estimated trend and cyclical components correspond to a HP lowpass filtered trend and a BK bandpass filtered cyclical component.

40. Harvey and Jaeger (1993) show that the HP trend can be replicated in a structural time series model by a smooth local linear trend with signal-to-noise ratio equal to the inverse of the HP smoothing parameter.

41. Ravn and Uhlig (2002).

References

Abowd, John and Arnold Zellner, 1985. "Estimating Gross Labor Force Flows." *Journal of Business and Economic Statistics* 3(3): 254–83.

Bartel, Ann P., 1979. "The Migration Decision: What Role Does Job Mobility Play?" *The American Economic Review* 69(5): 775–86.

Bureau of Labor Statistics, 2002. "Current Population Survey: Design and Methodology." Technical Paper 63RV, Washington, DC: U.S. Census Bureau.

Burns, Arthur F. and Wesley C. Mitchell, 1946. *Measuring Business Cycles*. National Bureau of Economic Research.

Chua, Tin Chiu and Wayne A. Fuller, 1987. "A Model for Multinomial Response Error Applied to Labor Flows." *Journal of the American Statistical Association* 82(397): 46–51.

Cleveland, William S., 1979. "Robust Locally Weighted Regression and Smoothing Scatterplots." *Journal of the American Statistical Association* 74(368): 829–36.

Dahmann, Donald C., 1986. "Geographical Mobility Research with Panel Data." *Growth and Change* 17(3): 35–48.

Davis, Steven J., R. Jason Faberman, John C. Haltiwanger, and Ian Rucker, 2008. "Adjusted Estimates of Worker Flows and Job Openings in JOLTS." Working Paper 14137, National Bureau of Economic Research.

Faberman, R. Jason, 2005. "Analyzing the JOLTS Hires and Separations Data." In "Proceedings of the Annual Statistical Association," Alexandria, VA: American Statistical Association.

Farber, Henry S., 1999. *Mobility and Stability: The Dynamics of Job Change in Labor Markets*, volume 3B of *Handbook of Labor Economics*, pp. 2439–83. Amsterdam: Elsevier Science, North-Holland.

Feng, Shuaizhang, 2001. "The Longitudinal Matching of Current Population Surveys: A Proposed Algorithm." *Journal of Economic and Social Measurement* 27(1–2): 71–91.

Fitzgerald, John, Peter Gottschalk, and Robert Moffitt, 1998. "An Analysis of Sample Attrition in Panel Data: The Michigan Panel Study of Income Dynamics." *The Journal of Human Resources* 33(2): 251–99.

Fujita, Shigeru and Garey Ramey, 2006. "The Cyclicality of Job Loss and Hiring." Working Paper 06-17, Federal Reserve Bank of Philadelphia.

Gallaway, Lowell E., 1969. "Age and Labor Mobility Patterns." *Southern Economic Journal* 36(2): 171–80.

Gottschalk, Peter and Robert Moffitt, 1999. "Changes in Job Instability and Insecurity Using Monthly Survey Data." *Journal of Labor Economics* 17(4): S91–S126.

Greenwood, Michael J., 1975. "Research on Internal Migration in the United States: A Survey." *Journal of Economic Literature* 13(2): 397–433.

Harvey, Andrew C., 1989. *Forecasting, Structual Time Series Models, and the Kalman Filter.* Cambridge University Press.

Harvey, Andrew C. and Albert Jaeger, 1993. "Detrending, Stylised Facts and the Business Cycle." *Journal of Applied Econometrics* 8(3): 231–47.

Harvey, Andrew C. and Thomas M. Trimbur, 2003. "General Model-Based Filters for Extracting Cycles and Trends in Economic Time Series." *The Review of Economics and Statistics* 85(2): 244–55.

Hausman, Jerry A., Jason Abrevaya, and Fiona M. Scott-Morton, 1998. "Misclassification of the Dependent Variable in A Discrete-Response Setting." *Journal of Econometrics* 87(2): 239–69.

Heckman, James J, 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47(1): 153–61.

Hodrick, Robert J. and Edward C. Prescott, 1997. "Post-war U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking* 29(1): 1–16.

Hogue, Carma R., 1985. "History of the Problems Encountered in Estimating Gross Flows." In "Proceedings of the Conference on Gross Flows in Labor Force Statistics," Washington, DC.

Hogue, Carma R. and Paul O. Flaim, 1986. "Measuring Gross Flows in the Labor Force: An Overview of a Special Conference." *Journal of Business and Economic Statistics* 4(1): 111–21.

Holzer, Harry J. and Robert J. LaLonde, 1999. "Job Change and Job Stability among Less-Skilled Young Workers." Discussion Paper 1191-99, Institute for Research on Poverty.

Jaeger, David A., 1997. “Reconciling the Old and New Census Bureau Education Questions: Recommendations for Researchers.” *Journal of Business and Economic Statistics* 12(3): 300–9.

Katz, Arnold, Klaus Teuter, and Philip Sidel, 1984. “Comparison of Alternative Ways of Deriving Panel Data from the Annual Demographic Files of the Current Population Survey.” *Review of Public Data Use* 12(1): 35–44.

Kletzer, Lori G., 1998. “Job Displacement.” *The Journal of Economic Perspectives* 12(1): 115–36.

Koopman, Siem Jan, Andrew C. Harvey, Jurgen A. Doornik, and Neil Shephard, 2007. *Structural Time Series Analyzer, Modeller and Predictor: STAMP 8*. London: Timberlake Consultants, Ltd.

Lansing, John B. and James N. Morgan, 1967. “The Effect of Geographical Mobility on Income.” *The Journal of Human Resources* 2(4): 449–60.

Madrian, Brigitte C. and John Lars Lefgren, 2000. “An Approach to Longitudinally Matching Current Population Survey (CPS) Respondents.” *Journal of Economic and Social Measurement* 26: 31âĂŞ–62.

Moscarini, Giuseppe and Kaj Thomsson, 2008. “Occupational and Job Mobility in the US.” *Scandinavian Journal of Economics* 109(4): 807–36.

Neal, Derek, 1999. “The Complexity of Job Mobility among Young Men.” *Journal of Labor Economics* 17(2): 237–261.

Nekarda, Christopher J., 2008. “Weekly Time Series of the U.S. Labor Market.” Unpublished paper, University of California, San Diego.

Neumark, David and Daiji Kawaguchi, 2004. “Attrition Bias in Labor Economics Research using Matched CPS Files.” *Journal of Economic and Social Measurement* 29(4): 445–72.

Peracchi, Franco and Finis Welch, 1994. “Trends in Labor Force Transitions of Older Men and Women.” *Journal of Labor Economics* 12(2): 210–42.

Poterba, James and Lawrence H. Summers, 1984. “Response Variation in the CPS: Caveats for the Unemployment Analysis.” *Monthly Labor Review* 107(3): 37–43.

Poterba, James and Lawrence H. Summers, 1986. “Reporting Errors and Labor Market Dynamics.” *Econometrica* 54(6): 1,319–38.

Ravn, Morten O. and Harald Uhlig, 2002. "On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations." *The Review of Economics and Statistics* 84(2): 371–80.

Schlottmann, Alan M. and Henry W. Herzog Jr., 1981. "Employment Status and the Decision to Migrate." *The Review of Economics and Statistics* 63(4): 590–98.

Schlottmann, Alan M. and Henry W. Herzog Jr., 1984. "Career and Geographic Mobility Interactions: Implications for the Age Selectivity of Migration." *The Journal of Human Resources* 19(1): 72–86.

Shimer, Robert, 2007. "Reassessing the Ins and Outs of Unemployment." Unpublished paper, University of Chicago.

Topel, Robert H. and Michael P. Ward, 1992. "Job Mobility and the Careers of Young Men." *The Quarterly Journal of Economics* 107(2): 439–479.

Tucker, C. Jack and William L. Urton, 1987. "Frequency of Geographic Mobility: Findings From the National Health Interview Survey." *Demography* 24(2): 265–70.

U.S. Census Bureau, 2007. "Annual Geographical Mobility Rates, By Type of Movement: 1947–2007." Internet release date 16 October 2007. Retrieved 16 November 2008 from http://www.census.gov/population/socdemo/migration/.

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

Table 1. Examples of Relationship among Addresses, Households, and Persons in the LPD[a]

	Interview history		
Example	*Address*	*Household*	*Person*
1. No noninterviews, no household change	iiii iiii	iiii iiii	iiii iiii
2. No noninterviews, household change	iiii iiii	iiii ____	iiii ____
	iiii iiii	____ iiii	____ iiii
3. Type C noninterview	iiii iC__	iiii i___	iiii i___
4. All noninterviews	BBBB BBC_	____ ____	____ ____
5. Intervening type B noninterview	iiiB iiii	iii_ ____	iii_ ____
	iiiB iiii	____ iiii	____ iiii
6. Trailing type B noninterview(s)	iiii iBBB	iiii i___	iiii i___
7. Initial type B noninterview(s)	BBBB iiii	____ iiii	____ iiii
8. An individual out-mover	iiii iiii	iiii iiii	iiii iiii
	iiii iiii	iiii iiii	iii_ ____
9. An individual in-mover	iiii iAii	iiii iAii	iiii iAii
	iiii iAii	iiii iAii	___i iAii

a. An i denotes a completed interview; A, B, and C denote type A, type B, and type C noninterviews; and a underscore denotes a missing observation. Each row per example depicts a separate household or person.

Table 2. Total Number of Longitudinal Units [a]

Unit	*Stayers*	*Movers*	*Total*
Addresses (AOU)	2,948,860	707,510	3,656,370
	[80.6]	[19.4]	[100.0]
Households (HOU)	2,959,414	1,201,421	4,160,835
	[71.1]	[28.9]	[100.0]
Persons [b] (POU)	7,984,300	2,617,606	10,601,906
	[75.3]	[24.7]	[100.0]

Source: Author's calculations using LPD data for 1976:1–2007:12.

a. Numbers in brackets are percent of total units.

b. Includes persons under sixteen years old.

Table 3. Number of Completed Interviews per POU, by Number of Attempted Interviews

Percent[a]

Attempted interviews	*Completed interviews*								
	1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*Total*
1	10.7								10.7
2	0.4	7.0							7.5
3	0.1	0.5	7.5						8.1
4	0.2	0.4	1.4	25.3					27.4
5	0.1	0.0	0.1	0.4	2.7				3.4
6	0.0	0.0	0.0	0.1	0.2	2.5			2.9
7	0.0	0.0	0.1	0.1	0.1	0.3	3.9		4.5
8	0.2	0.1	0.2	0.5	0.4	0.7	2.6	30.9	35.6
Total	11.7	8.1	9.3	26.5	3.4	3.5	6.5	30.9	100.0

Source: Author's calculations using LPD data for 1976:1–2007:12.

a. Share of validity-weighted count of person observation units (POUs) for persons aged sixteen years or older.

Table 4. Distribution of Demographic Characteristics, by Mobility

Percent[a]

Category	*LPD*	*Stayer*	*Out-mover*	*In-mover*
Sex				
Male	48.2	47.4	49.8	49.5
Female	51.9	52.6	50.2	50.5
Race				
White	84.3	85.7	82.2	81.1
Black	11.8	10.8	13.3	13.8
Other	4.0	3.5	4.6	5.1
Age				
16–24	21.2	15.5	31.9	32.9
24–54	54.9	55.0	53.2	55.4
55 and older	23.9	29.5	14.7	11.7
Marital status				
Married	57.6	63.2	45.7	46.8
Widowed or divorced	14.8	14.8	15.4	14.6
Never married	27.6	22.0	39.0	38.7
Education				
High school drop-out	23.9	23.6	24.6	24.6
High school graduate	35.8	35.8	35.2	35.9
Some college	21.3	21.0	22.4	21.8
College graduate	19.0	19.6	17.8	17.6

Source: Author's calculations using LPD data for 1976:1–2007:12.

a. Calculated using validity-weighted sampling weight for persons aged sixteen years or older.

Table 5. Distribution of Labor Force Characteristics, by Mobility

Percent [a]

Category	*LPD*	*Stayer*	*Out-mover*	*In-mover*
Labor force status				
E	61.8	61.2	63.4	64.6
U	4.0	3.3	6.3	6.8
N	34.2	35.6	30.3	28.7
Labor force transition				
EU	0.6	0.5	1.0	0.9
EN	1.3	1.2	1.6	1.3
UE	0.7	0.6	1.1	1.1
NE	1.1	1.1	1.2	1.2
UN	0.6	0.5	0.9	0.8
NU	0.6	0.5	0.9	0.8
XE	19.0	17.5	21.4	26.3
XU	1.4	1.0	2.3	3.2
XN	10.2	10.0	10.1	12.0

Source: Author's calculations using LPD data for 1976:1–2007:12.

a. Calculated using validity-weighted sampling weight for persons aged sixteen years or older.

Table 6. Effect of Geographic Mobility on Hazard Rates [a]

Hazard rate	*No. of obs.*	*Mean*	*Standard deviation*	*Minimum*	*Maximum*
Level [b]					
Separation	369	0.8048	0.0505	0.6593	0.9512
Job finding	369	0.9809	0.0335	0.8847	1.0779
Cyclical component [c]					
Separation	382	−0.0002	0.0143	−0.0382	0.0331
Job finding	382	0.0001	0.0037	−0.0075	0.0102

Source: Author's calculations using LPD data for 1976:1–2007:12.

a. Calculated using validity-weighted sampling weight for persons aged sixteen years or older.

b. Ratio of hazard rate calculated from population of stayers to that of entire population; see equation 7.

c. Difference between cyclical component estimated from population of stayers and that of entire population; see equation 9.

Table 7. Contemporaneous Cyclical Correlation with Unemployment[a]

Hazard rate	*All*	*Stayers*	*Bias*[b]
Separation	0.8743	0.8813	0.5537
Job finding	−0.9396	−0.9379	−0.1957

Source: Author's calculations using LPD data for 1976:1–2007:12.

a. corr $\left(\psi_t^{UR}, \psi_t^{y}\right)$, where y is the item listed in the column head.

b. Difference between cyclical component estimated from population of stayers and that of entire population; see equation 9.

Table 8. Separation and Job Finding Hazard Rates, Alternate Measures[a]

Ratio to validity-weighted measure

Measure	*No. of observations*	*Mean*	*Standard deviation*	*Minimum*	*Maximum*
Separation					
Naive	369	1.0000	0.0000	1.0000	1.0000
Threshold	366	0.9939	0.0081	0.9666	1.0182
Job finding					
Naive	369	1.0000	0.0000	1.0000	1.0000
Threshold	366	0.9995	0.0057	0.9755	1.0129

Source: Author's calculations using LPD data for 1976:1–2007:12.

a. Naive does not validate matches; threshold keeps all matches with $\overline{V}_i \geq 0.875$.

Table 9. Effect of Geographic Mobility on Hazard Rates, Alternate Measure[a]

Variable	No. of observations	Mean	Standard deviation	Minimum	Maximum
Level[b]					
Separation	382	0.8078	0.0524	0.6830	0.9684
Job finding	382	0.9820	0.0320	0.8642	1.0898
Cyclical component[c]					
Separation	382	0.0000	0.0522	−0.1391	0.1582
Job finding	382	0.0000	0.0316	−0.1203	0.1044

Source: Author's calculations using LPD data for 1976:1–2007:12.

a. Calculated using validity-weighted sampling weight for persons aged sixteen and older.

b. Ratio of hazard rate calculated from population of stayers to that of entire population; see equation 7.

c. Difference between cyclical component estimated from population of stayers and that of entire population; see equation 9. Cyclical component extracted from seasonally-adjusted series using the Hodrick-Prescott filter.

Table 10. Contemporaneous Cyclical Correlation with Unemployment, Alternate Measure[a]

Hazard rate	*All*	*Stayers*	*Bias*[b]
Separation	0.6042	0.5817	0.1180
Job finding	−0.7255	−0.6909	−0.0464

Source: Author's calculations using LPD data for 1976:1–2007:12.

a. corr $\left(\psi_t^{UR}, \psi_t^{y}\right)$, where y is the item listed in the column head. Cyclical component extracted from seasonally-adjusted series using the Hodrick-Prescott filter.

b. Difference between cyclical component estimated from population of stayers and that of entire population; see equation 9.

Table 11. LPD Variable Definitions

Variable name	*Description*	*Value*
AID	Unique address identifier	19-digit string
HID	Unique household identifier	20-digit string
PID	Unique person identifier	22-digit string
MONTH	Month of interview	2-digit number
YEAR	Year of interview	4-digit number
MIS	Month in sample	1-digit number [1-8]
INTSTAT	Interview status	1 – Interview 2 – Type A noninterview 3 – Type B noninterview 4 – Type C noninterview
STATE	FIPS state code	2-digit number [1-56]
SEX	Sex	1 – Male 2 – Female
RACE	Race	1 – White 2 – Black 3 – Other
AGE	Age at end of reference week	2-digit number [0-85], topcoded as 80 for ages 80–84 topcoded as 85 for ages 85+
MS	Marital status	1 – Married 2 – Widowed/divorced 3 – Never married
EDUC	Educational attainment	1 – Less than high school graduate 2 – High school graduate 3 – Some college 4 – College graduate
LFS	Labor force status	e – Employed u – Unemployed n – Not in the labor force

Table 11. LPD Variable Definitions (continued)

Variable name	*Description*	*Value*
IND	Major industry recode	2-digit number [1-23], see CPS data definitions
OCC	Major occupation recode	2-digit number [1-15], see CPS data definitions

Figure 1. CPS Sample Rotation Pattern

	Month																			
PID	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**
A	1	2	3	4									5	6	7	8				
B		1	2	3	4									5	6	7	8			
C			1	2	3	4									5	6	7	8		
D				1	2	3	4									5	6	7	8	
E					1	2	3	4									5	6	7	8
F						1	2	3	4									5	6	7
G							1	2	3	4									5	6
H								1	2	3	4									5
I									1	2	3	4								
J										1	2	3	4							
K											1	2	3	4						
L												1	2	3	4					
M													1	2	3	4				
N														1	2	3	4			
O															1	2	3	4		
P																1	2	3	4	
Q																	1	2	3	4

Figure 2. Hierarchy of Observation Units

Address Observation Unit

The collection of all interviews conducted at an address

→ Removed

AOUs with all type B or C noninterviews

+ Includes Type A and Type B noninterviews
+ Type C noninterviews end an AOU

↓

Household Observation Unit

The largest collection of observations within an AOU that could possibly come from the same household

→ Type B noninterviews between households

+ A Type B or Type C noninterview ends an HOU
+ Any remaining observations form a new HOU
+ A replacement household starts a new HOU

↓

Person Observation Unit

The largest collection of observations within a HOU that could possibly come from the same person

→ Nothing

+ Initial type A noninterviews remain with the person succeeding them
+ Intervening type A noninterviews do not start a new POU
+ Trailing type A noninterviews remain with the person preceding them

Figure 3. Distribution of Completed Interviews per POU, by Mobility

Percent[a]

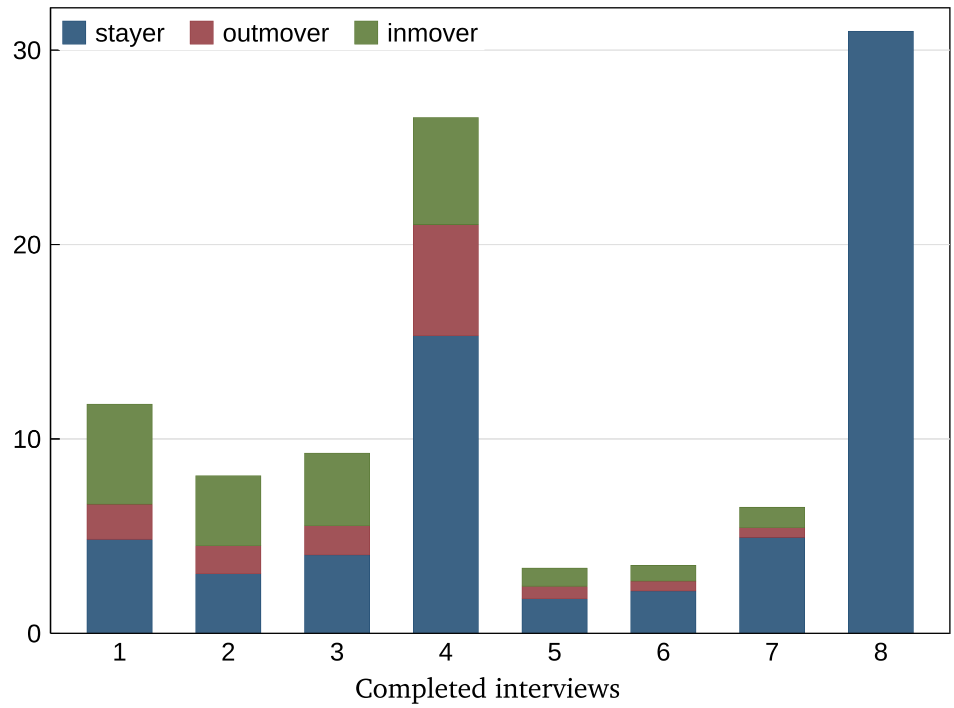

Source: Author's calculations using LPD microdata for 1976:1–2007:12.

a. Share of validity-weighted count of person observation units (POUs) for persons aged sixteen years or older.

Figure 4. Cyclical Behavior of Geographic Mobility, 1976–2007 [a]

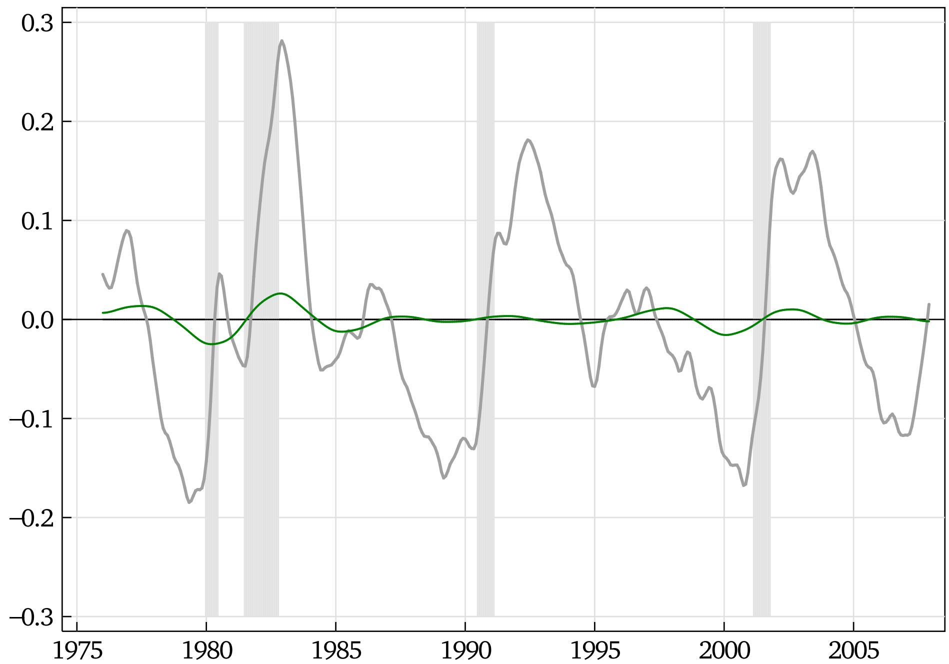

Source: Author’s calculations using annual data from U.S. Census Bureau (2007)

a. Cyclical components estimated using equation 8. Gray line is cyclical component of unemployment rate. Annual data are smoothed to monthly frequency using locally weighted polynomial regression smoother (Cleveland, 1979). Shaded regions indicate recessions as dated by the NBER.

Figure 5. Effect of Mobility on Cyclical Component of Separation Hazard Rate, 1976–2007[a]

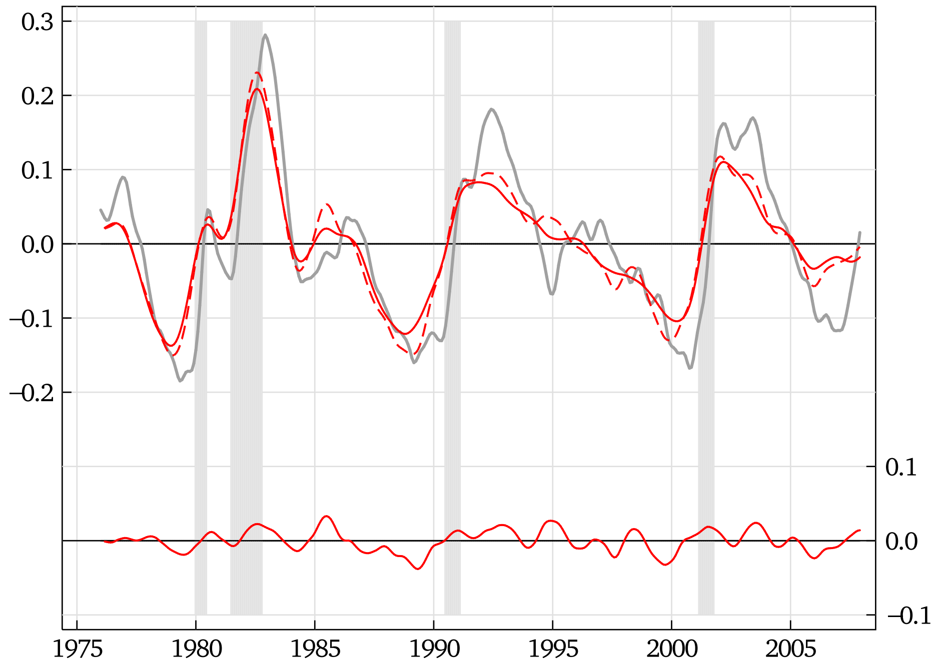

Source: Author's calculations using CPS microdata for 1976:2–2007:12.

a. Cyclical components estimated using equation 8. Gray line is cyclical component of unemployment rate. Thin solid line uses entire population; dashed line uses only population of stayers. Lower panel plots difference between cyclical component estimated from population of stayers and that of entire population; see equation 9. Shaded regions indicate recessions as dated by the NBER.

Figure 6. Effect of Mobility on Cyclical Component of Job Finding Hazard Rate, 1976–2007[a]

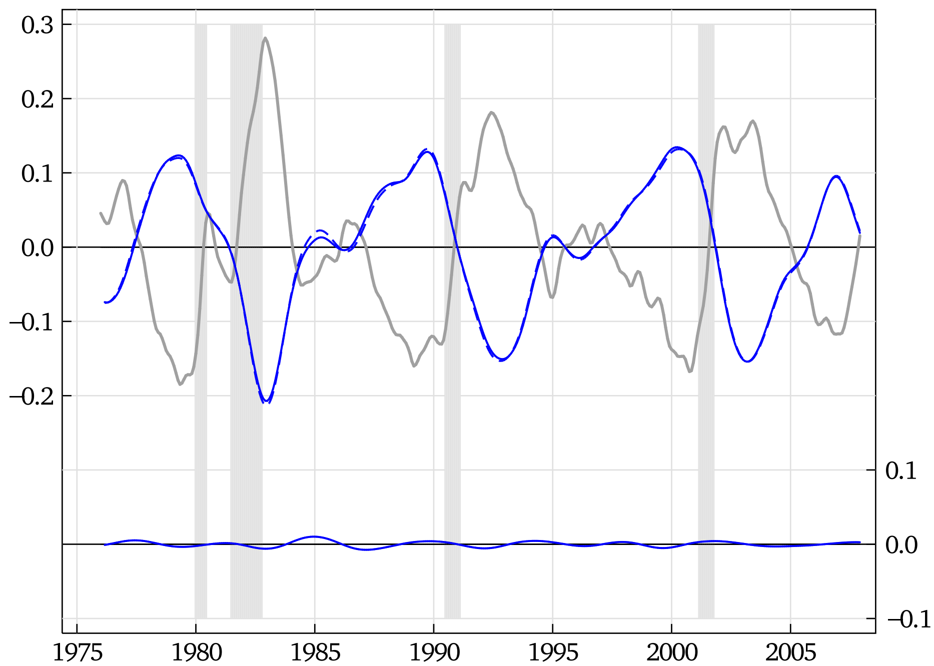

Source: Author's calculations using CPS microdata for 1976:2–2007:12.

a. Cyclical components estimated using equation 8. Gray line is cyclical component of unemployment rate. Thin solid line uses entire population; dashed line uses only population of stayers. Lower panel plots difference between cyclical component estimated from population of stayers and that of entire population; see equation 9. Shaded regions indicate recessions as dated by the NBER.

Figure 7. Effect of Mobility on Cyclical Component of Separation Hazard Rate, Alternate Measure, 1976–2007 [a]

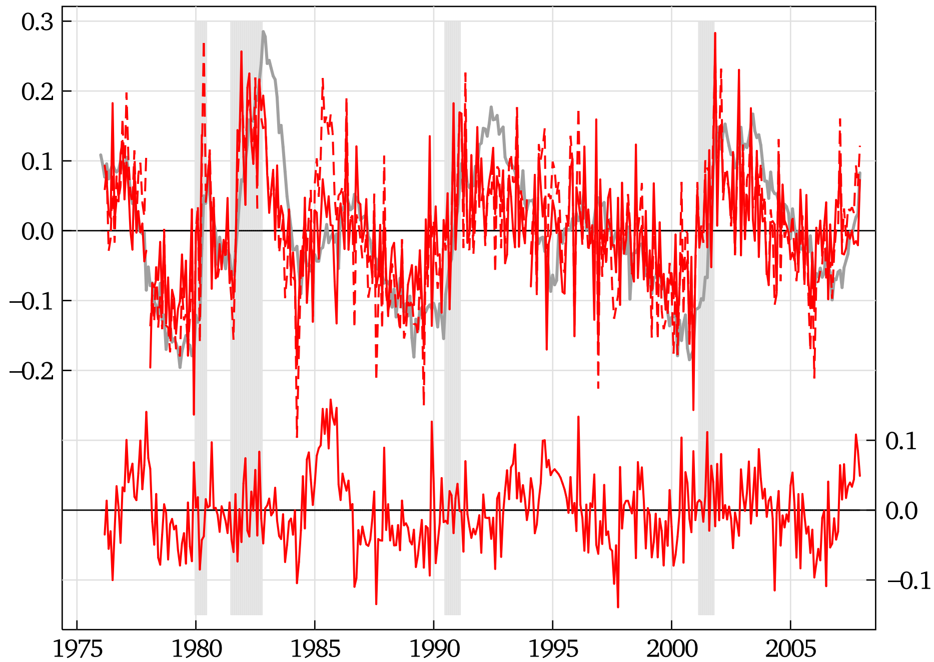

Source: Author's calculations using CPS microdata for 1976:2–2007:12.

a. Cyclical components estimated using equation 8. Gray line is cyclical component of unemployment rate. Solid maroon line uses entire population; dashed maroon line uses only population of stayers. Lower panel plots difference between cyclical component estimated from population of stayers and that of entire population; see equation 9. Shaded regions indicate recessions as dated by the NBER.

Figure 8. Effect of Mobility on Cyclical Component of Job Finding Hazard Rate, Alternate Measure, 1976–2007[a]

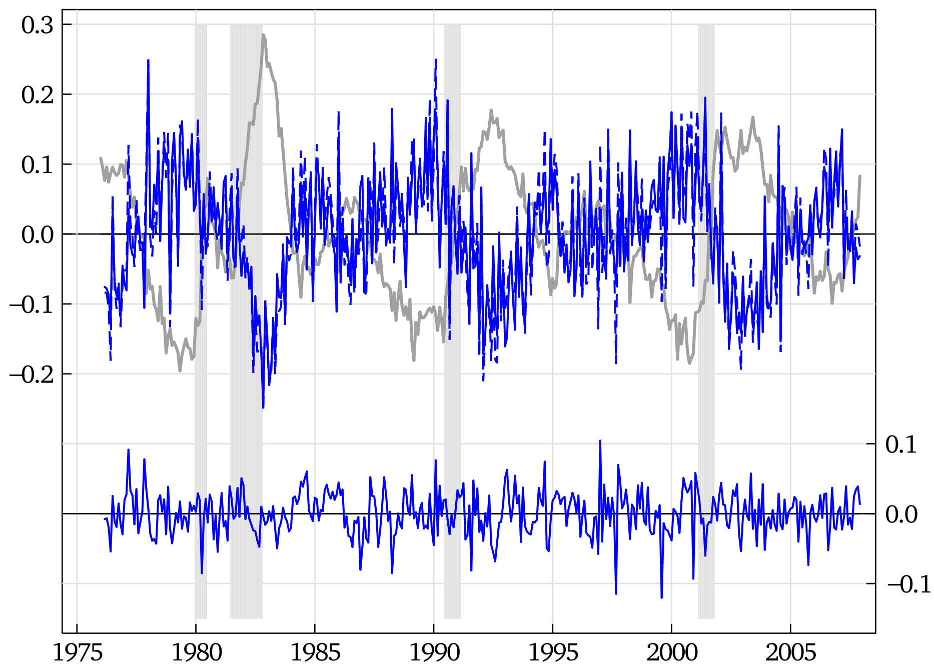

Source: Author's calculations using CPS microdata for 1976:2–2007:12.

a. Cyclical components estimated using equation 8. Dark gray line is cyclical component of unemployment rate. Solid navy line uses entire population; dashed navy line uses only population of stayers. Lower panel plots difference between cyclical component estimated from population of stayers and that of entire population; see equation 9. Shaded regions indicate recessions as dated by the NBER.